

Axes in Outer Space

Michael Handel and Lee Mosher

February 2, 2008

Abstract

We develop a notion of axis in the Culler–Vogtmann outer space \mathcal{X}_r of a finite rank free group F_r , with respect to the action of a nongeometric, fully irreducible outer automorphism ϕ . Unlike the situation of a loxodromic isometry acting on hyperbolic space, or a pseudo-Anosov mapping class acting on Teichmüller space, \mathcal{X}_r has no natural metric, and ϕ seems not to have a single natural axis. Instead our axes for ϕ , while not unique, fit into an “axis bundle” \mathcal{A}_ϕ with nice topological properties: \mathcal{A}_ϕ is a closed subset of \mathcal{X}_r proper homotopy equivalent to a line, it is invariant under ϕ , the two ends of \mathcal{A}_ϕ limit on the repeller and attractor of the source–sink action of ϕ on compactified outer space, and \mathcal{A}_ϕ depends naturally on the repeller and attractor. We propose various definitions for \mathcal{A}_ϕ , each motivated in different ways by train track theory or by properties of axes in Teichmüller space, and we prove their equivalence.

Contents

1	Introduction	3
1.1	Characterizations of the axis bundle.	4
1.2	The main theorems.	8
1.3	A question of Vogtmann.	9
1.4	Contents and proofs.	9
1.5	Problems and questions	12
2	Preliminaries	15
2.1	Outer space and outer automorphisms	15
2.2	Paths, Circuits and Edge Paths	20
2.3	Folds	21
2.4	Train Track Maps	24
2.5	The attracting tree T_+	28
2.6	Geodesic laminations in trees and marked graphs	33
2.7	The expanding lamination Λ_-	35
2.8	Relating Λ_- to T_- and to T_+	38

3	The ideal Whitehead graph	40
3.1	Definition and structure of the ideal Whitehead graph	41
3.2	Asymptotic leaves and the ideal Whitehead graph	43
3.3	T_+ and the ideal Whitehead graph	44
3.4	An example of an ideal Whitehead graph	45
4	Cutting and pasting local stable Whitehead graphs	48
4.1	Pasting local stable Whitehead graphs	48
4.2	Cutting local stable Whitehead graphs	50
4.3	The finest local decomposition	51
5	Weak train tracks	53
5.1	Local decomposition of the ideal Whitehead graph.	54
5.2	Folding up to a weak train track	55
5.3	Comparing train tracks to weak train tracks	57
5.4	Rigidity and irrigidity of Λ_- isometries	59
5.5	Examples of exceptional weak train tracks	62
6	Topology of the axis bundle	66
6.1	Continuity properties of the normalized axis bundle	66
6.2	The Gromov topology on weak train tracks	68
6.3	Properness of the length map.	71
6.4	Applying Skora's method to the Properness Theorem 6.1	74
7	Fold lines	82
7.1	Examples of fold paths	82
7.2	Characterizing fold lines	84
7.3	Direct limits of fold rays	86
7.4	Legal laminations of split rays	89
7.5	Weak train tracks on fold lines	93

1 Introduction

There are many interesting and fruitful analogies between the group of isometries of hyperbolic n -space \mathbf{H}^n , the mapping class group of a finite type surface acting on Teichmüller space, and the outer automorphism group $\text{Out}(F_r)$ of a rank r free group F_r acting on outer space \mathcal{X}_r .

In all cases the space has a natural compactification, obtained by adding a boundary at infinity to which the group action extends continuously. In all cases there is a class of “hyperbolic” elements whose extended action has source–sink dynamics, with a repeller and an attractor at infinity: loxodromic isometries of \mathbf{H}^n ; pseudo-Anosov mapping classes of a surface [McC85]; and outer automorphisms of F_r which are *fully irreducible* [LL03], meaning that no proper, nontrivial free factor of F_r has a conjugacy class which is periodic under the action of the outer automorphism.

In \mathbf{H}^n and in Teichmüller space, a “hyperbolic” element ϕ has a unique axis in the usual metric sense, a properly embedded, oriented, ϕ -invariant geodesic line which ϕ translates in the forward direction. The axis of ϕ depends only on the ordered pair (repeller, attractor), and this dependence is natural with respect to the group action: the axis is the unique geodesic line whose negative end converges to the repeller and whose positive end converges to the attractor.

In \mathcal{X}_r , however, it is not reasonable to expect such a nice notion of an axis for a fully irreducible $\phi \in \text{Out}(F_r)$. There does not seem to be a natural candidate for a metric on \mathcal{X}_r , and certainly not one that has the unique geodesic properties of \mathbf{H}^n or Teichmüller space. So among the many properly embedded lines in \mathcal{X}_r whose two ends converge to the repeller and the attractor of ϕ , it seems hard to single out one of those lines in a natural way and call it *the* axis of ϕ . Nevertheless, among all such lines we shall select in a natural manner a particular collection of them, the union of which we call the “axis bundle” of ϕ , a subset $\mathcal{A}_\phi \subset \mathcal{X}_r$ with the following properties:

Properness: \mathcal{A}_ϕ is a closed subset of \mathcal{X}_r proper homotopy equivalent to \mathbf{R} . The negative end of \mathcal{A}_ϕ limits in compactified outer space on the repeller, and the positive end limits on the attractor.

Naturality: \mathcal{A}_ϕ depends only on the ordered pair (repeller, attractor) in a manner which is natural with respect to the action of $\text{Out}(F_r)$.

We shall do this only in the case where ϕ is *nongeometric*, meaning that it does not arise from a homeomorphism of a compact surface with boundary. The geometric case, while conceptually much simpler and more well understood (see [BH92]), has some peculiarities whose inclusion in our theory would overburden an already well laden paper.

Naturality means that for any nongeometric, fully irreducible $\phi, \phi' \in \text{Out}(F_r)$ and any $\psi \in \text{Out}(F_r)$, if ψ takes the repeller–attractor pair of ϕ to the repeller–attractor pair of ϕ' , then ψ takes \mathcal{A}_ϕ to $\mathcal{A}_{\phi'}$. One consequence of naturality is that \mathcal{A}_ϕ is invariant under ϕ , because the repeller–attractor pair of ϕ is invariant under ϕ . Another is that axes are indifferent to iterates (positive powers) of ϕ : ϕ^n and ϕ have the same axis for all $n > 0$, because ϕ^n has the same repeller and the same attractor as ϕ .

Certain differences between axes in hyperbolic space or Teichmüller space on the one hand, and axis bundles in outer space on the other hand, may already be apparent from

the above features. First, the topological relation between the axis bundle and the real line is not homeomorphism but instead proper homotopy equivalence. Second, axes in \mathbf{H}^n or Teichmüller space are indifferent to inverses — the axis of ϕ^{-1} is the orientation reversal of the axis of ϕ . But as we have argued in [HM06b], it seems unnatural to insist that axes in \mathcal{X}_r be indifferent to inverses, due to the phenomenon that a fully irreducible $\phi \in \text{Out}(F_r)$ can have an expansion factor which is different than the expansion factor of its inverse. Although our paper [HM06a] offers a bare hint of an “indifference to inverses” property, by bounding the ratio of the logarithms of the expansion factors of ϕ and ϕ^{-1} , nevertheless we do not attempt here to construct a notion of axis which is indifferent to inverses.

The main results of this work are a theorem which proves equivalence of several characterizations of the axis bundle — from some of which the naturality feature of the axis bundle is evident — and a theorem which proves the properness feature of the axis bundle. In the remainder of this introduction we shall describe the three characterizations of the axis bundle, and the methods of proof.

1.1 Characterizations of the axis bundle.

We shall characterize the axis bundle by presenting three different definitions, the equivalence of which is one of our major theorems. These definitions are motivated by various aspects of the theory of outer automorphisms, and by analogies with various properties of Teichmüller geodesics.

Fold lines. Fold lines form a natural class of paths in outer space arising from several concepts: train track maps due to Bestvina and Handel [BH92]; Stallings folds [Sta83]; and semiflow concepts behind Skora’s method for investigating homotopy properties of spaces of \mathbf{R} -trees ([Sko], and see also [Whi93], [Cla05], [GL05]).

Recall that elements of compactified outer space $\overline{\mathcal{X}}_r = \mathcal{X}_r \cup \partial\mathcal{X}_r$ are represented by minimal, very small actions of F_r on \mathbf{R} -trees, what we shall call “ F_r -trees”. Two F_r -trees represent the same element of $\overline{\mathcal{X}}_r$ if and only if they are related by a homothety (a bijection that multiplies the metric by a constant) that is F_r -equivariant. An element of \mathcal{X}_r itself is represented by a simplicial F_r -tree on which the action of F_r is free. Passing to the quotient modulo F_r of a free, simplicial F_r -tree, an element of outer space is also represented by a “marked graph”, a metric graph G without valence 1 vertices whose fundamental group is marked by an isomorphism $F_r \approx \pi_1(G)$; while we allow G to have valence 2 vertices, their presence or absence does not effect the element of \mathcal{X}_r that G represents. The group $\text{Out}(F_r)$ acts on $\overline{\mathcal{X}}_r$ from the right: the image under $\phi \in \text{Out}(F_r)$ of a marked graph G is a marked graph $G \cdot \phi$ defined by precomposing the marking $F_r \approx \pi_1(G)$ with an automorphism of F_r that represents ϕ .

Given a fully irreducible $\phi \in \mathcal{X}_r$, the repeller and attractor for ϕ are elements of $\partial\mathcal{X}_r$ represented by F_r -trees that we denote $T_- = T_-^\phi$ and $T_+ = T_+^\phi$.

A *fold line* in outer space is a continuous, injective, proper function $\mathbf{R} \mapsto \mathcal{X}_r$ defined by a continuous 1-parameter family of marked graphs $t \mapsto G_t$ for which there exists a family of homotopy equivalences $h_{ts}: G_s \rightarrow G_t$ defined for $s \leq t \in \mathbf{R}$, each of which preserves marking, such that the following hold:

Train track property: h_{ts} is a local isometry on each edge for all $s \leq t \in \mathbf{R}$.

Semiflow property: $h_{ut} \circ h_{ts} = h_{us}$ for all $s \leq t \leq u \in \mathbf{R}$, and $h_{ss}: G_s \rightarrow G_s$ is the identity for all $s \in \mathbf{R}$.

Our first definition of the axis bundle is:

Definition (1): Fold lines. \mathcal{A}_ϕ is the union of the images of all fold lines $\gamma: \mathbf{R} \rightarrow \mathcal{X}_r$ such that $\gamma(t)$ converges in $\overline{\mathcal{X}}_r$ to T_-^ϕ as $t \rightarrow -\infty$ and to T_+^ϕ as $t \rightarrow +\infty$.

Notice that, just from the format of this definition, the axis bundle \mathcal{A}_ϕ has the desired naturality feature, because the collection of fold lines in \mathcal{X}_r is invariant under the action of $\text{Out}(F_r)$, and the points $T_-^\phi, T_+^\phi \in \partial \mathcal{X}_r$ depend naturally on ϕ .

The construction of ϕ -periodic fold lines is given in detail in the beginning of Section 7, drawing on well known ideas reviewed in Section 2.2. Here is an outline of the construction.

Recall that every fully irreducible $\phi \in \text{Out}(F_r)$ is represented by an *affine train track map*, a homotopy equivalence $g: G \rightarrow G$ of a marked graph such that g takes vertices to vertices, g induces the outer automorphism ϕ on $\mathcal{F}_r \approx \pi_1(G)$, and there exists $\lambda > 1$ so that the restriction of each iterate g^i to the interior of each edge locally stretches the metric by the factor λ^i ([BH92], and see Section 2.4). The number $\lambda = \lambda(\phi)$ is independent of g and is called the *expansion factor* of ϕ . Let $G_i = \frac{1}{\lambda^i} G \cdot \phi^i$ be renormalized marked graphs representing the orbit of G under powers of ϕ . For integers $i \leq j$ the iterate $g^{j-i}: G \rightarrow G$ of the map g induces a marking preserving homotopy equivalence $g_{ji}: G_i \rightarrow G_j$, and these maps satisfy the semiflow property, which is just a fancy way of saying that the following diagram commutes:

$$\begin{array}{ccccccc} \cdots & \longrightarrow & G_i & \xrightarrow{g_{i+1,i}} & G_{i+1} & \longrightarrow & \cdots \longrightarrow G_{j-1} \xrightarrow{g_{j,j-1}} G_j \longrightarrow \cdots \\ & & & \searrow & & & \nearrow \\ & & & & g_{ji} & & \end{array}$$

Moreover, the definition of train track maps and the renormalization of G_i together imply that g_{ji} satisfies the above train track property. Now one uses Stallings method of folds to factor the map $g_{1,0}: G_0 \rightarrow G_1$ as a product of fold maps, and then one interpolates each fold map continuously to obtain a path in \mathcal{X}_r from G_0 to G_1 . Concatenating the ϕ^i translates of this path for all $i \in \mathbf{Z}$, one obtains a ϕ -periodic fold line in \mathcal{X}_r . Source-sink dynamics on compact subsets of outer space [BFH97] implies that in $\overline{\mathcal{X}}_r$ this fold line converges in the negative direction to T_- and in the positive direction to T_+ . This shows that the point of \mathcal{X}_r represented by G satisfies Definition (1) of \mathcal{A}_ϕ .

Notice that we do *not* claim that \mathcal{A}_ϕ is the union of all ϕ -periodic fold lines. In fact, we do not even claim that \mathcal{A}_ϕ is the union of all ϕ^i -periodic fold lines over all $i \geq 1$. See Section 5.3 for an example of $G \in \mathcal{A}_\phi$ that does lie on any ϕ^i -periodic fold line for any iterate ϕ^i . In particular, such a G is not the domain of a train track representative of any iterate ϕ^i , nor is G the domain of a *nonsimplicial train track representative*¹ of any iterate, a map that satisfies the definition of a train track representative except that it is not required to take vertices to vertices.

¹also called an “efficient representative” in an unpublished preprint of T. White and in [LL04].

An analogy with Teichmüller geodesics. Fold lines can also be motivated by an analogy with Teichmüller geodesics, as follows. Recall that if S is a closed oriented surface of genus ≥ 2 , each oriented geodesic γ in the Teichmüller space $\mathcal{T}(S)$ has the following essentially unique description. Choosing a unit speed parameterization $t \mapsto \sigma_t$ of γ , there is a transverse pair of measured foliations $\mathcal{F}_-, \mathcal{F}_+$ with transverse measures denoted $|dy|, |dx|$, respectively, such that for each $t \in \mathbf{R}$ the singular Euclidean metric $ds_t = \sqrt{(e^t |dx|)^2 + (e^{-t} |dy|)^2}$ agrees with the conformal structure σ_t . If $|dx|, |dy|$ are normalized so that $\int_S |dx| \cdot |dy| = 1$ then the pair $\mathcal{F}_-, \mathcal{F}_+$ depends up to joint isotopy only on the parameterization of γ . Also, in Thurston's boundary $\partial\mathcal{T}(S) = \mathcal{PMF}(S)$, the projective measured foliation space, the pair of points that is represented by $\mathcal{F}_-, \mathcal{F}_+$ depends only on γ as an oriented geodesic in $\mathcal{T}(S)$. Notice that the identity map on S underlies a system of maps $\sigma_s \mapsto \sigma_t$ that satisfies the semiflow property. Consider now the renormalized metric $ds'_t = e^{-t} ds_t = \sqrt{(|dx|)^2 + (e^{-2t} |dy|)^2}$, also in agreement with the conformal structure σ_t . For each leaf L of \mathcal{F}_- and each $t \leq u \in \mathbf{R}$, the restrictions to L of the metric ds'_t and the metric ds'_u are in agreement, because each agrees with the restriction of the transverse measure $|dx|$. To put it another way, each of the maps $\sigma_t \mapsto \sigma_u$ restricts to a local isometry on each leaf L of \mathcal{F}_- , with respect to the metrics ds'_t and ds'_u . Thus we have demonstrated an analogy with the maps $G_s \mapsto G_t$ in the definition of a fold line, each of which restricts to a local isometry on each edge of G_s .

An even stronger analogy between Teichmüller geodesics and fold lines will be used below to motivate Definition (2).

Laminations, trees, and Λ_- isometries. Thurston compactified Teichmüller space with projective classes of measured laminations [FLP⁺79]. Skora proved a duality theorem between \mathbf{R} -trees and measured foliations, which provides an alternative description of the boundary of Teichmüller space in terms of \mathbf{R} -trees [Sko90]. The boundary of outer space $\partial\mathcal{X}_r$ is described solely in terms of \mathbf{R} -trees, no general theory of measured laminations having so far arisen.

However, in [BFH97] Bestvina, Feighn, and Handel define the *expanding lamination* of a fully irreducible $\phi \in \text{Out}(F_r)$ in terms of train track representatives of ϕ , and they prove a duality relation between the expanding lamination and T_- which is a natural analogue of Skora's duality. This relation is described by saying that the expanding lamination of ϕ is the “length zero lamination” of T_- (see Section 3 of [BFH97], and Lemma 2.19 below). Because of this duality we shall denote the expanding lamination of ϕ as $\Lambda_- = \Lambda_-^\phi$. The lamination Λ_- can be regarded abstractly as an F_r -equivariant subset of $\partial F_r \times \partial F_r$, or it may be realized as a family of leaves — bi-infinite geodesics — in any free, simplicial F_r -tree T ; see Sections 2.6 and 2.7 below.

Given a free, simplicial F_r -tree T , an F_r -equivariant map $f_T: T \rightarrow T_+$ is called a Λ_- *isometry* if, for each leaf L of Λ_- realized in T , the restriction of f_T to L is an isometry onto a bi-infinite geodesic in T_+ .

Here is our next definition of the axis bundle:

Definition (2): Supports of Λ_- isometries to T_+ . \mathcal{A}_ϕ is the set of all elements of \mathcal{X}_r represented by a free simplicial F_r tree T for which there exists a Λ_- isometry $f_T: T \rightarrow T_+$.

For example, given any train track representative $g: G \rightarrow G$ of ϕ , the universal cover $T = \tilde{G}$, regarded as a free, simplicial F_r -tree, satisfies Definition (2). The details of why this is true are given in Corollary 2.20, but here is the idea. By lifting to the universal cover T_i of each marked graph G_i in the diagram above, we obtain F_r -equivariant maps $\tilde{g}_{ji}: T_i \rightarrow T_j$ that satisfy the semiflow property for integer values $i \leq j$, that is, the following diagram commutes for all $i \leq j$:

$$\begin{array}{ccccccc} \cdots & \longrightarrow & T_i & \xrightarrow{\tilde{g}_{i+1,i}} & T_{i+1} & \longrightarrow & \cdots \longrightarrow T_{j-1} \xrightarrow{\tilde{g}_{j,j-1}} T_j \longrightarrow \cdots \\ & & & \searrow & & & \nearrow \\ & & & & \tilde{g}_{ji} & & \end{array}$$

Each map \tilde{g}_{ji} is an F_r -equivariant lift of the train track map $g^{j-i}: G \rightarrow G$, and the restriction of \tilde{g}_{ji} to each edge of T_i is an isometry. In this situation, the trees T_i and the maps \tilde{g}_{ji} form a direct system whose direct limit is an F_r tree representing T_+ , and from the direct limit construction one obtains a map $f_i: T_i \rightarrow T_+$ for each i . Moreover, from the construction of the expanding lamination Λ_- it follows that for each leaf L of Λ_- realized in each T_i , the restriction of g_{ji} to L is an isometry. Passing to the direct limit, it follows that the restriction of f_i to L is an isometry, and so the map f_i demonstrates that the point of \mathcal{X}_r represented by T_i or by G_i satisfies Definition (2) of \mathcal{A}_ϕ .

Another analogy with Teichmüller geodesics. Further motivation for Definition (2) comes again by an analogy with geodesics in Teichmüller space $\mathcal{T}(S)$. Consider such a geodesic γ with all the associated data as above: $\sigma_t, \mathcal{F}_-, \mathcal{F}_+, |dy|, |dx|, ds_t, ds'_t$. Lifting to the universal cover \tilde{S} , the measured foliation $\tilde{\mathcal{F}}_+$ with transverse measure $|dx|$ is dual to an \mathbf{R} -tree T_+ on which $\pi_1(S)$ acts: as a point set T_+ is obtained from \tilde{S} by collapsing each leaf of $\tilde{\mathcal{F}}_+$ to a point; as a metric space T_+ is obtained by thinking of $|dx|$ as a pseudo-metric on \tilde{S} and passing to the associated metric space obtained by collapsing to a point each maximal subset of pseudo-metric diameter zero. Let $f: \tilde{S} \rightarrow T_+$ be the collapse map. For each t , the conformal structure σ_t lifts to a $\pi_1(S)$ equivariant conformal structure $\tilde{\sigma}_t$, and we write the map f in the form $f_t: \tilde{\sigma}_t \rightarrow T_+$. Under these conditions, the restriction of f_t to each leaf L of $\tilde{\mathcal{F}}_-$ realized in the singular Euclidean structure ds'_t is an isometry, because the metric ds'_t restricted to L equals the transverse measure $|dx|$ restricted to L . Thus we can say that the map f_t is an \mathcal{F}_- -isometry (with respect to the metric ds'_t), and so each point σ_t on γ satisfies a property that is an analogue of Definition (2). To state this property precisely, given a geodesic γ determined by the data above, we say that $\sigma \in \mathcal{T}(S)$ satisfies property (2') if and only if there exists a singular Euclidean structure ds agreeing with σ , and there is a map $f: \tilde{S} \rightarrow T_+$, such that locally f is modelled on the projection of a Euclidean rectangle to one of its sides, and for each leaf L of \mathcal{F}_- realized as a \tilde{ds} geodesic in \tilde{S} , the restriction of f to L is an isometry onto a bi-infinite geodesic in T_+ . We have shown that each point on γ satisfies (2'), and the converse is easily seen to hold as well. Points of γ are thus characterized by property (2'), just as points in the axis bundle \mathcal{A}_ϕ are characterized by (2).

Train tracks. Here is our final definition of the axis bundle, which could have been stated first but has been relegated to the last. Given a fully irreducible $\phi \in \text{Out}(F_r)$, while “train track representatives” of ϕ are well established objects of study, “train tracks” themselves

are not. We define a *train track* for ϕ to be a point of \mathcal{X}_r that is represented by a marked graph G such that there exists an affine train track representative $g: G \rightarrow G$ of ϕ . Let $\text{TT}(\phi) \subset \mathcal{X}_r$ be the set of train tracks of ϕ .

Definition (3): Closure of train tracks. \mathcal{A}_ϕ is the closure of $\bigcup_{i=1}^{\infty} \text{TT}(\phi^i)$ in \mathcal{X}_r .

Remark. In Section 5.5 we produce an element of \mathcal{A}_ϕ which is not a train track. We also believe there are easy examples showing that the closure of $\text{TT}(\phi)$ itself need not equal \mathcal{A}_ϕ , although we have not produced such an example. This raises the question: does there exist some $N \geq 1$ such that \mathcal{A}_ϕ is the closure of $\bigcup_{i=1}^N \text{TT}(\phi^i)$?

1.2 The main theorems.

Having proposed three definitions of the axis bundle, our first main result shows that these three definitions are equivalent:

Theorem 1.1. *For each nongeometric, fully irreducible $\phi \in \text{Out}(F_r)$ with repelling and attracting points $T_-, T_+ \in \partial\mathcal{X}_r$ and length zero laminations Λ_-, Λ_+ , definitions (1), (2), and (3) of \mathcal{A}_ϕ are equivalent. In other words, for any $G \in \mathcal{X}_r$ the following are equivalent:*

- (1) *G lies on some fold line $t \mapsto G_t$, $t \in \mathbf{R}$, such that $G_t \rightarrow T_-$ as $t \rightarrow -\infty$ and $G_t \rightarrow T_+$ as $t \rightarrow +\infty$.*
- (2) *There exists a Λ_- isometry from the universal cover of G to T_+ .*
- (3) *G is in the closure of the set of train tracks of iterates of ϕ .*

The set \mathcal{A}_ϕ defined by this theorem is called the *axis bundle* of ϕ . While Definitions (1) and (3) may seem the most intuitive, Definition (2) occupies an important middle ground between intuition and rigor, being easier to work with when proving certain properties, in particular properness. For much of the paper we focus on Definition (2), and in Section 5 we introduce the terminology “weak train track” to refer to G satisfying Definition (2). Of course, once Theorem 1.1 is proved, then a weak train track is any element of the axis bundle \mathcal{A}_ϕ .

Note particularly that the natural dependence of \mathcal{A}_ϕ on the ordered pair $(T_-, T_+) \in \partial\mathcal{X}_r \times \partial\mathcal{X}_r$ is an immediate consequence of definition (1) and also of definition (2).

Our next main result establishes the properness feature of the axis bundle. Recall that a proper map $f: A \rightarrow B$ is a *proper homotopy equivalence* if there are proper maps $\bar{f}: B \rightarrow A$, $H: A \times [0, 1] \rightarrow B$, and $\bar{H}: B \times [0, 1] \rightarrow A$ such that H is a homotopy between $\bar{f} \circ f$ and Id_A , and \bar{H} is a homotopy between $f \circ \bar{f}$ and Id_B .

Theorem 1.2. *If $\phi \in \text{Out}(F_r)$ is nongeometric and fully irreducible, and if $\gamma: \mathbf{R} \rightarrow \mathcal{X}_r$ is a fold line such that $\gamma(t) \rightarrow T_\pm^\phi$ as $t \rightarrow \pm\infty$, then the inclusion map $\gamma(\mathbf{R}) \hookrightarrow \mathcal{A}_\phi$ is a proper homotopy equivalence. Moreover, the end of \mathcal{A}_ϕ corresponding to the negative end of L converges in $\overline{\mathcal{X}_r}$ to T_- , and the end of \mathcal{A}_ϕ corresponding to the positive end of L converges in $\overline{\mathcal{X}_r}$ to T_+ .*

1.3 A question of Vogtmann.

In [Vog02] Vogtmann asked:

“Is the set of train tracks for an irreducible automorphism contractible?”

We address this question only for the case of a nongeometric, fully irreducible $\phi \in \text{Out}(F_r)$. The question requires some interpretation, because the notion of “train tracks” for ϕ is not well established in the literature. The notion of train tracks that we have defined above — points in outer space that support affine, irreducible train track representatives of ϕ — is not appropriate, because the set of train tracks of ϕ has only countably many orbits under the action of the cyclic group $\langle \phi \rangle$ on \mathcal{X}_r .

By interpreting Vogtmann’s question instead using weak train tracks — elements of the axis bundle of ϕ — Theorem 1.2 gives a strong positive answer to her question, establishing not just the homotopy type but the property homotopy type of the set of weak train tracks A_ϕ .

Relations to work of Los and Lustig. Another approach to Vogtmann’s question is given in a preprint of Los and Lustig [LL04], who define a space of nonsimplicial train track representatives which they prove is contractible. One of our key technical results, Proposition 5.4, is very similar to Lemma 3.16 of [LL04] which also plays a key role in their results. Our results differ from theirs in several aspects. They study a space of maps between marked graphs representing elements of \mathcal{X}_r , whereas our axis bundle is a subset of \mathcal{X}_r ; see also the remarks at the end of Section 6.2. Their proof uses an interesting semiflow which is necessarily different from the semiflow used by Skora in [Sko], whereas we prove proper homotopy equivalence to the line by plugging directly into Skora’s theorem; see Section 6.4. Finally, in Section 5.5 we give an example of a weak train track that is not the domain of a nonsimplicial train track representative; hence, although the subset of \mathcal{X}_r consisting of the domains of nonsimplicial train track representative can be shown to be a dense, contractible subset of the axis bundle, this subset is not equal to the whole axis bundle, and so it is not closed in \mathcal{X}_r , and its proper homotopy type is not evident.

1.4 Contents and proofs.

First we describe the contents of the individual sections, then we prove the two main theorems by referring to results found in various sections.

Section 2, Preliminaries. We give very detailed preliminaries, some of which is a compilation of previous work. Some details give new geometric perspectives of familiar objects such as train track representatives (Section 2.4) and laminations (Sections 2.6 and 2.7). Other details arise in making explicit what was only implicit in previous works, and in providing detailed proofs of results that, while previously evident, have not previously been recorded (see in particular Sections 2.5, Section 2.7, and Section 2.8).

Section 3, The ideal Whitehead graph. In this section and the next we introduce the study of the *ideal Whitehead graph* of a nongeometric, fully irreducible $\phi \in \text{Out}(F_r)$. Intuitively this graph describes topological and combinatorial aspects of the singularities

of ϕ . The graph has one edge for each singular leaf of the expanding lamination of ϕ , with two singular half-leaves impinging on the same vertex whenever those half-leaves are asymptotic. The analogous construction carried out for a singularity of the unstable foliation of a pseudo-Anosov surface homeomorphism produces a circle with three or more vertices. In $\text{Out}(F_r)$ the topology of a component of the ideal Whitehead graph can be different from a circle, as we show by an example in Section 3.4.

The ideal Whitehead graph of ϕ is defined in Section 3.1 in terms of the actions of automorphisms representing ϕ on the boundary of F_r . In terms of a train track representative $g: G \rightarrow G$, we show that the ideal Whitehead graph of ϕ decomposes into the local stable Whitehead graphs of g , in a pattern dictated by the Nielsen paths of g . Sections 3.2 and 3.3 give other descriptions of the ideal Whitehead graph, in terms of: asymptotic relations among leaves of Λ_- ; and branch points of the attracting F_r -tree T_+ .

Section 4, Cutting and pasting local stable Whitehead graphs. Given a nongeometric, fully irreducible $\phi \in \text{Out}(F_r)$ and a train track representative $g: G \rightarrow G$, we study how the decomposition of the ideal Whitehead graph of ϕ into local stable Whitehead graphs of g varies under some standard methods of varying g that were first described in [BH92]. This variation takes the form of cutting and pasting local stable Whitehead graphs, according to the formation or collapse of Nielsen paths of g . The main result of this section, Proposition 4.4, describes in a natural way those train track representatives which give the finest possible decomposition of the ideal Whitehead graph.

Section 5, Weak train tracks. In this section we take up in earnest the study of the axis bundle, by proving several results about weak train tracks. We show that many structural features of train tracks, e.g. the induced local decomposition of the ideal Whitehead graph, have analogues for weak train tracks.

A key technical result, Proposition 5.4, says roughly that if one fixes a train track G then for every weak train track H , one can obtain H from G by a sequence of Stallings folds provided two conditions hold: the local decomposition of H is no finer than the local decomposition of G ; and the length of H is sufficiently small in terms of the geometry of G . As mentioned above, this result is closely related to a key result of [LL04].

Another important result in this section, Proposition 5.5, shows how to tell when a weak train track is a train track. As a corollary we obtain one of the implications of Theorem 1.1 by proving that a dense set of weak train tracks are train tracks.

Also, in Section 5.4 we characterize when a weak train track H is the domain of a unique Λ_- isometry, and more generally we describe the space of all Λ_- isometries with domain H .

Section 6, Topology of the axis bundle. This section proves several of our central results.

Theorem 6.1 describes a proper homotopy equivalence, the *length map*, from the set of weak train tracks to the positive real axis $(0, \infty)$. This theorem represents most of the work in proving the Properness Theorem 1.2. The proof is an application of Skora's semiflow methods [Sko], which have also played a central role in other investigations of the homotopy types of spaces of \mathbf{R} -trees [Whi93], [Cla05], [GL05]. We shall directly apply the main theorem of [Sko], for which we must set up the appropriate maps and verify the hypotheses of

that theorem. This is accomplished by applying the main results of the earlier sections, in particular Proposition 5.4 and the study of rigidity and irrigidity carried out in Section 5.4.

Also proved in this section is the fact that the set of weak train tracks is a closed subset of outer space, a step in the proof of Theorem 1.1 (see below).

Section 7, Fold Lines. In this section we initiate a detailed study of fold lines. The main result proved is the equivalence of Definition (1) of the axis bundle with the other definitions of the axis bundle. The hard work here is the proof that definition (1) implies definition (2): every point on a fold line is a weak train track. To do this, given a fold line $t \mapsto G_t$ in \mathcal{X}_r , using only the assumption that $\lim_{t \rightarrow \pm\infty} G_t = T_{\pm}$ in $\overline{\mathcal{X}}_r$ we show that both the expanding lamination Λ_- and the attracting tree T_+ can be reconstructed by natural geometric processes stated in terms of the marked graphs G_t and the edge isometries $h_{ts}: G_s \rightarrow G_t$. These processes allow us to construct a Λ_- isometry from each G_t to T_+ , proving that G_t is a weak train track.

Along the way we also prove one other important piece in the proof of the Properness Theorem 1.2, that the length map restricted to each fold line is a homeomorphism between that fold line and the positive real axis.

Throughout the paper we have probably erred on the side of more detail rather than less, particularly when the extra detail seems structurally interesting and can be illustrated with particular examples of train track maps; see for example Sections 3.4 and 5.4. We have tried to develop the theory of the axis bundle comprehensively, incorporating three different natural definitions into the statement of Theorem 1.1. However, we have isolated all the material regarding Definition (1) in Section 7, and so the reader who is willing to completely ignore Definition (1), and consider only those portions of the theory concerned with Definitions (2) and (3) will find the paper somewhat shortened.

The Proof of Theorem 1.1. Let $\phi \in \text{Out}(F_r)$ be nongeometric and fully irreducible, and denote $T_{\pm} = T_{\pm}^{\phi}$, $\Lambda_{\pm} = \Lambda_{\pm}^{\phi}$. We use the terminology *weak train track* of ϕ to refer to an element of \mathcal{X}_r satisfying Definition (2), being represented by a free, simplicial F_r -tree T for which there exists a Λ_- isometry $T \mapsto T_+$.

The equivalence of definitions (1), (2), and (3) breaks into several implications.

Definition (2) implies Definition (3): A dense set of weak train tracks are train tracks, as proved in Corollary 5.6.

Definition (3) implies Definition (2): Every train track for ϕ is a weak train track, as proved in Lemma 2.20. The set of weak train tracks is closed, as proved in Lemma 6.2 (1). Together these prove that the closure of the set of train tracks is contained in the set of weak train tracks.

Definition (2) implies Definition (1): Every weak train track of ϕ lies on a fold line for ϕ , as proved by a construction given in Proposition 7.1.

Definition (1) implies Definition (2): Every point on a fold line for ϕ is a weak train track for ϕ , as proved in Theorem 7.2, which is stated in Section 7, and whose proof takes up the bulk of Section 7. \diamond

The Proof of Theorem 1.2. Let \mathcal{A} be the axis bundle of a nongeometric, fully irreducible $\phi \in \text{Out}(F_r)$. In the beginning of Section 6 we define the *length map* $\text{Len}: \mathcal{A} \rightarrow (0, +\infty)$,

by normalizing each weak train track in a natural manner and taking the length of this normalization, and we prove that the length map satisfies the equivariance condition

$$\text{Len}(T \cdot \phi^i) = \lambda(\phi)^{-i} \text{Len}(T) \quad \text{for any } i \in \mathbf{Z}$$

Theorem 6.1 says that $\text{Len}: \mathcal{A} \rightarrow (0, +\infty)$ is a proper homotopy equivalence.

Consider a fold line $\gamma: \mathbf{R} \rightarrow \mathcal{X}_r$ such that $\gamma(t)$ converges in $\overline{\mathcal{X}}_r$ to T_-^ϕ as $t \rightarrow -\infty$ and to T_+^ϕ as $t \rightarrow +\infty$. From Theorem 1.1 it follows that $\gamma(\mathbf{R}) \subset \mathcal{A}$. Corollary 7.3 of Theorem 6.1 says that the restriction of the length map $\text{Len}: \mathcal{A} \rightarrow (0, +\infty)$ to $\gamma(\mathbf{R})$ is a homeomorphism. It follows that the injection $\gamma(\mathbf{R}) \hookrightarrow \mathcal{A}$ is a proper homotopy equivalence.

By combining the equivariance condition with the fact that $\text{Len}: \mathcal{A} \rightarrow (0, +\infty)$ is a proper homotopy equivalence, it follows that the set $D = \text{Len}^{-1}[1, \lambda(\phi)]$ is a compact fundamental domain for the action of the cyclic group $\langle \phi \rangle$ on \mathcal{A} . Moreover, the end of \mathcal{A} that corresponds to the positive end of the fold line $\gamma(\mathbf{R})$ is represented by the subset $\cup_{i=1}^{+\infty} \phi^i(D)$. The unique limit point of this set in $\overline{\mathcal{X}}_r$ is T_+ , because of the result from [BFH00] that the action of ϕ^i on $\overline{\mathcal{X}}_r$ converges uniformly on compact subsets of \mathcal{X}_r to the constant map with value T_+ , as $i \rightarrow +\infty$. Similarly, the end of \mathcal{A} that corresponds to the negative end of $\gamma(\mathbf{R})$ converges to T_- . \diamond

1.5 Problems and questions

Combinatorial structure of the axis bundle. Outer space has a combinatorial structure, a subdivision into closed subsets called “cells”, each parameterized by a simplex minus some subcomplex of that simplex, so that the intersection of any two cells is a common face; see [CV86] or Section 2.1 below.

Question 1. Given a nongeometric, fully irreducible $\phi \in \text{Out}(F_r)$, do the intersections of \mathcal{A}_ϕ with the cells of outer space define a cell decomposition of \mathcal{A}_ϕ ?

The Properness Theorem 1.2 provides a proper homotopy equivalence from \mathcal{A}_ϕ to $(0, \infty)$ called the *length function*. If the previous question has an affirmative answer, perhaps level sets of the length function can be studied using the Morse theory ideas of Bestvina and Brady [BB97], in order to address the following:

Question 2. Are the level sets of the length function contractible?

Geometry of the axis bundles of ϕ and of ϕ^{-1} . The Properness Theorem 1.2 implies that \mathcal{A}_ϕ has two ends, and that there exist subsets of \mathcal{A}_ϕ that are compact and that separate the two ends; we call such a subset a *cross section* of \mathcal{A}_ϕ . Our proof of the Properness Theorem carried out in Section 6 provides a construction of a cross section, using a Stallings fold argument. Fold ideas are also central to the techniques of [HM06a] relating the expansion factors of ϕ and of ϕ^{-1} , and we believe that these techniques may extend to relating their axis bundles as well. We propose here a question that would serve to explore this relation.

To quantify the size of a cross section, fix an appropriate measurement of (pseudo)distance on \mathcal{X}_r , for example: an $\text{Out}(F_r)$ -equivariant geodesic metric for which closed balls are compact; or a combinatorial measurement where the distance between two points $x, y \in \mathcal{X}_r$ is

the minimal length of a chain of cells $\sigma_0, \sigma_1, \dots, \sigma_L$ with $x \in \sigma_0$, $y \in \sigma_L$, and $\sigma_{i-1} \cap \sigma_i \neq \emptyset$ for all $i = 1, \dots, L$. Given a point $x \in \mathcal{A}_\phi$ define the *girth of \mathcal{A}_ϕ at x* to be the infimal diameter of a cross section of \mathcal{A}_ϕ passing through x , and define the *girth of \mathcal{A}_ϕ* to be the supremal girth over all points $x \in \mathcal{A}_\phi$. Given $x \in \mathcal{A}_\phi \cup \mathcal{A}_{\phi^{-1}}$, define the *joint girth at x of \mathcal{A}_ϕ and $\mathcal{A}_{\phi^{-1}}$* to be the infimal diameter of a compact subset of \mathcal{X}_r which contains the point x and which contains cross-sections of \mathcal{A}_ϕ and $\mathcal{A}_{\phi^{-1}}$. Define the *joint girth of \mathcal{A}_ϕ* to be the supremum over all $x \in \mathcal{A}_\phi \cup \mathcal{A}_{\phi^{-1}}$ of the joint girth at x .

There is no uniform bound for girth depending only on rank. The example in Section 7.1 of a nonconvergent split ray can be extended to a fold line for some reducible element $\phi \in \text{Out}(F_r)$, and with some thought one sees that the union of fold lines for ϕ contains a plane properly embedded into \mathcal{X}_r . By carefully interjecting a few more folds into the fold description of ϕ , one can produce nongeometric, fully irreducible elements of $\text{Out}(F_r)$ with axis bundles of arbitrarily large girth.

Given a nongeometric, fully irreducible $\phi \in \text{Out}(F_r)$, the ideas of [HM06a] give a correlation between the expansion factors of ϕ and of ϕ^{-1} , and they do so in a manner which, we believe, also gives a correlation between large girth of \mathcal{A}_ϕ and some kind of quantitative measurement that tells us “how irreducible” ϕ is. We would therefore expect to obtain some control over the joint girths of the axis bundles of ϕ and ϕ^{-1} , as long as one of them is “very irreducible”.

Problem 3. Fixing the rank r , give conditions on a nongeometric, fully irreducible $\phi \in \text{Out}(F_r)$ which imply a bound on the girth of \mathcal{A}_ϕ . Also, give conditions that would imply a bound on the joint girth of \mathcal{A}_ϕ and $\mathcal{A}_{\phi^{-1}}$.

Generalizations of axes. In an attempt to find a broader generalization of Teichmüller geodesics, it would be interesting to obtain results not just about source sink pairs (T_-, T_+) for fully irreducible outer automorphisms, but for more arbitrary pairs of points in $\partial\mathcal{X}_r$.

Question 4. For what pairs of points (T_0, T_1) in $\partial\mathcal{X}_r$ is there a good theory of axes, mimicking our results for source sink pairs (T_-, T_+) ?

Singularity types. As mentioned above, in Section 3 we present the ideal Whitehead graph of a nongeometric, fully irreducible outer automorphism as a description of the singularity structure of that outer automorphism.

First recall the singularity structure of a pseudo-Anosov homeomorphism $f: S \rightarrow S$ of a closed, oriented surface S . The *index* of a singularity s of f is defined to be $i(s) = 1 - \frac{k}{2}$ where s has k unstable directions and k stable directions. The Euler-Poincaré formula gives an index equation $\sum_s i(s) = \chi(S)$, summed over all singularities of f .

Consider now a nongeometric, fully irreducible $\phi \in \text{Out}(F_r)$. If C is a component of the ideal Whitehead graph, define its *index* to be the number

$$i(C) = 1 - \frac{\#\text{vertices}(C)}{2}$$

The *index* of ϕ is defined to be $i(\phi) = \sum_C i(C)$ summed over all components of the ideal Whitehead graph (modulo the action of F_r). In contrast to the situation on a surface, rather than an index equation there is in general only an index inequality $i(\phi) \geq 1 - r$: in [GJLL98]

the F_r -tree index $i(T_+^\phi)$ is defined and is proved to satisfy $i(T_+^\phi) \geq 1 - r$, and the results of Section 3 imply that $i(\phi) = i(T_+^\phi)$. In [BF] it is proved that the inequality $i(\phi) \geq 1 - r$ is an equality if and only if T_+^ϕ is a geometric F_r -tree, which in our present context occurs if and only if ϕ is a parageometric outer automorphism.

Question 5. What values of the index deficit $i(\phi) - (1 - r)$ are possible? Does the maximum index deficit go to $+\infty$ as $r \rightarrow +\infty$?

As a finer invariant, define the *index type* of ϕ to be the list of indices of the components of the ideal Whitehead graph of ϕ , written in increasing order. For instance when $r = 3$ then the possible index types which sum to $1 - r = -2$ are $(-\frac{1}{2}, -\frac{1}{2}, -\frac{1}{2}, -\frac{1}{2})$, $(-\frac{1}{2}, -\frac{1}{2}, -1)$, $(-\frac{1}{2}, -\frac{3}{2})$, and $(-1, -1)$, and the possible index types whose sum is strictly greater than -2 are $(-\frac{1}{2}, -\frac{1}{2}, -\frac{1}{2})$, $(-\frac{1}{2}, -\frac{1}{2})$, $(-\frac{1}{2})$, and (-1) . See Section 3.4 for an example in F_3 where the index type is (-1) . Masur and Smillie proved that on a finite type surface, all possible index types, which sum to the Euler characteristic, are achieved by pseudo-Anosov homeomorphisms, with a few exceptions in low genus [MS93].

Question 6. What possible index types, which satisfy the index inequality $i(\phi) \geq 1 - r$, are achieved by nongeometric, fully irreducible elements of $\text{Out}(F_r)$?

Even finer than the index type is the isomorphism type of the ideal Whitehead graph of ϕ . In the analogous situation of a pseudo-Anosov surface homeomorphism $f: S \rightarrow S$, a singularity s of f has an ideal Whitehead graph whose isomorphism type is determined by the index: the unstable directions at s are arranged in a circle, and so if the index equals i then the ideal Whitehead graph of s is just a circle with $-1 - 2i$ vertices and edges. In contrast, the example given in Section 3.4 is of a nongeometric, fully irreducible $\phi \in \text{Out}(F_3)$ having a connected ideal Whitehead graph homeomorphic to the triod, more precisely having 5 vertices and 4 edges in the topology of a tree with one vertex of valence 3, three of valence 1, and one of valence 2.

Question 7. What isomorphism types of graphs actually occur among the ideal Whitehead graphs of fully irreducible outer automorphisms of F_r ?

2 Preliminaries

2.1 Outer space and outer automorphisms

Outer automorphisms. Once and for all fix the integer $r \geq 2$. Let $F_r = \langle x_1, \dots, x_r \rangle$ be the free group of rank r with free basis x_1, \dots, x_r . Let $\text{Aut}(F_r)$ be the automorphism group. Let $i_\gamma: F_r \rightarrow F_r$ denote the inner automorphism $i_\gamma(\delta) = \gamma\delta\gamma^{-1}$. Let $\text{Inn}(F_r) = \{i_\gamma: \gamma \in F_r\}$ be the inner automorphism group. Let

$$\text{Out}(F_r) = \text{Aut}(F_r)/\text{Inn}(F_r)$$

be the outer automorphism group, which acts on conjugacy classes of elements and of subgroups of F_r .

An outer automorphism $\phi \in \text{Out}(F_r)$ is *reducible* if there exists a nontrivial free decomposition $F_r = A_1 * \dots * A_r * B$ such that ϕ permutes the conjugacy classes of A_1, \dots, A_r . If ϕ is not reducible then it is *irreducible*. If ϕ^i is irreducible for all integers $i \geq 1$ then ϕ is *fully irreducible*; this is also known in the literature as “IWIP”, an acronym for “irreducible with irreducible powers”. ϕ is *geometric* if there exists a homeomorphism $f: S \rightarrow S$ of a compact surface with boundary S , and an isomorphism $F_r \approx \pi_1(S)$ which conjugates ϕ the outer automorphism of $\pi_1(S)$ induced by f . If ϕ is not geometric then it is *nongeometric*.

Fix a *standard rose* R_r , a graph with one vertex v and directed edges e_1, \dots, e_r , and identify $F_r = \pi_1(R_r, v)$ where x_i is the class of the loop e_i . This induces an isomorphism of $\text{Aut}(F_r)$ with the group of homotopy classes of homotopy equivalences of the pair (R_r, v) , and an isomorphism of $\text{Out}(F_r)$ with the group of homotopy classes of homotopy equivalences of R_r . We use these isomorphisms without comment henceforth.

Outer space. A *graph* is a finite 1-complex G in which every vertex v has $\text{valence}(v) \geq 2$. The *rank* of G is the rank of $\pi_1 G$. The valence condition implies that every homotopy equivalence between graphs is surjective. In a graph of rank r , an Euler characteristic argument shows that $\sum_v \frac{1}{2}(\text{valence}(v) - 2) = r - 1$, and so there are at most $2r - 2$ vertices of valence ≥ 3 . Also, $\text{valence}(v) \leq 2r$ for each v , and if in addition $\text{valence}(v) \geq 3$ for each vertex v then G has at most $3r - 3$ edges.

A *metric graph* is a graph G equipped with a geodesic metric. A *similarity* or *homothety* between two metric graphs G, G' is a homeomorphism $h: G \rightarrow G'$ such that $d(hx, hy) = \lambda d(x, y)$ for some constant $\lambda > 0$ called the *stretch factor* of h . We use the notation $\text{Len}(\cdot)$ to represent length of objects in a metric graph G . For example: $\text{Len} G$ is the sum of the length of the edges of G ; for a path $p: [a, b] \rightarrow G$ we use $\text{Len} p$ for the path length; and if c represents a homotopy class of maps $S^1 \rightarrow G$ then we use $\text{Len} c$ to represent the minimal length of a representative of c . Sometimes to clarify the context we use $\text{Len}_G(\cdot)$ to mean length in the graph G .

A *marked graph* is a metric graph G equipped with a homotopy equivalence, usually denoted $\rho = \rho_G: R_r \rightarrow G$, called a *marking* of G . Each marking ρ_G induces an isomorphism $F_r \approx \pi_1(G, \rho_G(v))$, and this sets up a bijection between the set of markings of G modulo homotopy and the set of isomorphisms $F_r \approx \pi_1(G)$ modulo inner automorphism; we use this correspondence henceforth without comment. Given two marked graphs G, G' , a homotopy equivalence $f: G \rightarrow G'$ *preserves marking* if $f \circ \rho_G$ is homotopic to $\rho_{G'}$. Two marked graphs G, G' are *(isometrically) equivalent* if there exists an isometry $G \mapsto G'$ that preserves

marking, and G, G' are *projectively equivalent* if there exists a similarity $G \mapsto G'$ that preserves marking.

Culler and Vogtmann's *outer space* \mathcal{X}_r is the set of projective equivalence classes of marked graphs [CV86]. Define $\widehat{\mathcal{X}}_r$ to be the set of (isometric) equivalence classes of marked graphs. By our conventions, a marked graph may contain valence 2 vertices, and may be subdivided or unsubdivided at such points without altering the equivalence class in \mathcal{X}_r or in $\widehat{\mathcal{X}}_r$. With this flexibility, our definition of \mathcal{X}_r is equivalent to that of [CV86] where a marked graph has no valence 2 vertices.

We shall review below the topologies on \mathcal{X}_r and $\widehat{\mathcal{X}}_r$, after discussing the interpretation of these spaces in terms of **R**-trees.

Notation: Often we blur the distinction between a marked graph G , its isometric equivalence class in $\widehat{\mathcal{X}}_r$, and even its projective equivalence class in \mathcal{X}_r . When we need more formality, we use (ρ_G, G) as notation for a marked graph, and $[\rho_G, G]$ for its equivalence class, projective or nonprojective as the context makes clear.

The group $\text{Out}(F_r)$ acts on the right of \mathcal{X}_r and $\widehat{\mathcal{X}}_r$ by changing the marking: given a marked graph (ρ_G, G) , and given $\phi \in \text{Out}(F_r)$ corresponding to a homotopy equivalence $h: R_r \rightarrow R_r$, the action is defined by

$$[\rho_G, G] \cdot \phi = [\rho_G \circ h, G]$$

Any homotopy equivalence $g: G \rightarrow G$ of a marked graph (ρ_G, G) determines a homotopy equivalence $h = \bar{\rho}_G \circ g \circ \rho_G: R_r \rightarrow R_r$ well-defined up to homotopy independent of the choices of ρ_G and its homotopy inverse $\bar{\rho}_G$ within their homotopy classes, and so g determines an outer automorphism ϕ of F_r ; we say that g *represents* ϕ . We will always assume that g is an immersion on edges. If in addition g takes vertices to vertices then we say that g is a *topological representative* of ϕ . Note that

$$[\rho_G, G] \cdot \phi = [\rho_G \circ h, G] = [g \circ \rho_G, G]$$

Actions on trees. We review group actions on **R**-trees, focussing on the group F_r .

An **R**-tree is a metric space T such that for any $x \neq y \in T$ there exists a unique embedded topological arc, denoted $[x, y] \subset T$, with endpoints x, y , and this arc is isometric to the interval $[0, d(x, y)] \subset \mathbf{R}$. Given $x \in T$, two arcs $[x, y], [x, y']$ determine the *same germ* or the *same direction* at x if $[x, y] \cap [x, y'] \neq \{x\}$. This is an equivalence relation on arcs with endpoint x , and the set of equivalence classes is called the set of *germs* or *directions* at x . A *triad* based at a point $x \in T$ is a union of three arcs with endpoints at x no two of which have the same germ at x . The *valence* of x is the cardinality of the set of directions, equivalently, the number of components of $T - x$. The valence is ≥ 3 if and only if there exists a triad based at x . We say that T is *simplicial* if the points of valence $\neq 2$ form a discrete set. If T is simplicial and valence is everywhere finite then T is homeomorphic to a CW-complex.

Every isometry $f: T \rightarrow T$ of an **R**-tree is either *elliptic* meaning that f has a fixed point, or *hyperbolic* meaning that f has an *axis* ℓ , an f -invariant isometrically embedded copy of the real line on which the *translation length* $\inf_{x \in T} d(x, f(x))$ is minimized. In the elliptic case the translation length equals zero.

Consider an isometric action of F_r on an \mathbf{R} -tree T , denoted $\mathcal{A}: F_r \rightarrow \text{Isom}(T)$. Given $\gamma \in F_r$ and $x \in T$ we use various notations including t_γ for $\mathcal{A}(\gamma)$ and $\gamma \cdot x$ for $t_\gamma(x) = \mathcal{A}(\gamma)(x)$. For a subset $S \subset T$ we denote $\gamma \cdot S = \{\gamma \cdot x \mid x \in S\}$. The action \mathcal{A} is *minimal* if there does not exist a nonempty, proper, F_r -invariant subtree; this implies that T has no valence 1 points, because the valence 1 points are F_r -invariant and their complement is a nonempty F_r -invariant subtree. The action \mathcal{A} is *very small* if the stabilizer of every nondegenerate arc is either trivial or a cyclic group generated by a primitive element of F_r , and the stabilizer of every triod is trivial. We shall refer to a very small, minimal action of F_r on an \mathbf{R} -tree T as an F_r -tree.

To each F_r -tree T we associate the *length function* $\ell_T \in \mathbf{R}^C$, where $\ell_T[\gamma]$ is the translation length of the action of t_γ on T .

Consider two F_r -trees T, T' . A map $f: T \rightarrow T'$ is *equivariant* if $\gamma \cdot f(x) = f(\gamma \cdot x)$ for $x \in T$. We say that T, T' are *equivariantly isometric* or *isometrically conjugate* if there is an equivariant isometry $T \rightarrow T'$, and they are *topologically conjugate* if there exists an equivariant homeomorphism. A result of Culler and Morgan [CM87] says that T, T' are isometrically conjugate if and only if $\ell_T = \ell_{T'}$ in \mathbf{R}^C .

Outer space in terms of trees. Given a marked graph (ρ_G, G) with universal covering map $pr: \tilde{G} \rightarrow G$, we always assume that a basepoint in the universal cover \tilde{G} has been chosen which projects to $\rho_G(v)$, thereby inducing an action of F_r on \tilde{G} by covering translations. This makes \tilde{G} a free, simplicial F_r -tree whose isometric conjugacy class is well-defined independent of choices. Conversely, every free, simplicial F_r -tree T determines a quotient marked graph $G_T = T/F_r$, whose marking has homotopy class depending only on the isometric conjugacy class of T . Denoting $\tilde{G} = T_{G_T}$, clearly we have $T_{G_T} = T$ and $G_{T_{G_T}} = G$, and so we obtain a natural bijection between $\hat{\mathcal{X}}_r$ and the isometric conjugacy classes of free, simplicial F_r -trees. This bijection agrees with the length map to \mathbf{R}^C : if $T = \tilde{G}$ then $\ell_T = \ell_G$ in \mathbf{R}^C . Under this bijection the right action of $\text{Out}(F_r)$ on $\hat{\mathcal{X}}_r$ can be described as follows: given a free, simplicial F_r tree T with action $\mathcal{A}: F_r \rightarrow \text{Isom}(T)$, and given $\phi \in \text{Out}(F_r)$, choose $\Phi \in \text{Aut}(F_r)$ representing ϕ , and let $T \cdot \phi$ be the F_r tree with underlying \mathbf{R} -tree T and with action $\mathcal{A} \circ \Phi: F_r \rightarrow \text{Isom}(T)$; this action is well-defined up to conjugacy independent of the choice of Φ .

We can therefore describe $\hat{\mathcal{X}}_r$ *either* as marked graphs up to marked isometry *or* as free, simplicial F_r -trees up to isometric conjugacy; outer space \mathcal{X}_r has a similar pair of descriptions. Given a free, simplicial F_r -tree T , depending on context we use $[T]$ for the point in \mathcal{X}_r or $\hat{\mathcal{X}}_r$ represented by T , but we very often abuse notation by deleting the brackets and letting T stand for the F_r -tree as well as the corresponding point in \mathcal{X}_r or $\hat{\mathcal{X}}_r$.

Ends, automorphisms, and lifts. The Gromov boundary of the free group F_r is a Cantor set denoted ∂F_r . The action of $\text{Aut}(F_r)$ on F_r extends to an action on ∂F_r ; the extension of $\Phi \in \text{Aut}(F_r)$ is denoted $\hat{\Phi}: \partial F_r \rightarrow \partial F_r$.

Given a free simplicial F_r -tree T , the space of ends of T is denoted ∂T , and there is a natural homeomorphic identification $\partial F_r \approx \partial T$ which we will use without further comment. Given $\gamma \in F_r$, the map $t_\gamma: T \rightarrow T$ extends to a map of ends denoted $\hat{t}_\gamma: \partial T \rightarrow \partial T$, whose fixed point set is equal to the pair of endpoints t_γ^\pm of the axis of t_γ .

Suppose that G is marked graph and $g: G \rightarrow G$ is a homotopy equivalence representing $\phi \in \text{Out}(F_r)$. A lift of g to the F_r -tree \tilde{G} will be denoted $\breve{g}: \tilde{G} \rightarrow \tilde{G}$, where the “\breve” notation is meant to be a reminder that the lifted map is *not* F_r -equivariant (unless ϕ is the identity and g preserves marking). The map \breve{g} extends to a map of ends denoted $\hat{g}: \partial\tilde{G} \rightarrow \partial\tilde{G}$. There is a bijection between the set of lifts $\breve{g}: \tilde{G} \rightarrow \tilde{G}$ of g and the set of automorphisms $\Phi: F_r \rightarrow F_r$ representing ϕ . This bijection is characterized by the *twisted equivariance equation*

$$\breve{g} \circ t_\gamma = t_{\Phi(\gamma)} \circ \breve{g} \quad \text{for all } \gamma \in F_r. \quad (2.1)$$

The bijection is also characterized by the condition that \breve{g} and Φ extend to the same homeomorphism $\hat{g} = \hat{\Phi}: \partial F_r \rightarrow \partial F_r$. A simple computation using Equation 2.1 shows that if $\breve{g}, \breve{g}': \tilde{G} \rightarrow \tilde{G}$ are the lifts of g with corresponding automorphisms $\Phi, \Phi' \in \text{Aut}(F_r)$ that represent ϕ , and if $\gamma \in F_r$ is the unique element such that $\Phi' = i_\gamma \circ \Phi$, then $\breve{g}' = t_\gamma \circ \breve{g}$.

Notation: The notations \hat{g} and $\hat{\Phi}$ are used interchangeably, depending on context.

The topology of outer space. The topology on \mathcal{X}_r has three equivalent descriptions which we review: the cellular topology; the length or axes topology; and the Gromov topology. We use all three of these descriptions at various places in the paper.

The cellular topology [CV86]. Consider a marked graph G . With respect to the natural vertex set consisting of points of valence ≥ 3 , denote the edge set by $E(G)$. The *open cone* of G is the set of elements in $\hat{\mathcal{X}}_r$ represented by marked graphs G' for which there exists a homeomorphism $G \mapsto G'$ that preserves marking; the *open cell* of G in \mathcal{X}_r is similarly defined. Note that if $g: G \rightarrow G$ is a homeomorphism of a marked graph that preserves marking then g takes each edge to itself. It follows that if G' is in the open cone of $\hat{\mathcal{X}}_r$ containing G then any two homeomorphisms $G \mapsto G'$ that preserve marking induce the same bijection $E(G) \leftrightarrow E(G')$, and therefore the lengths of edges on G' determine a well-defined vector $L(G') \in [0, \infty)^{E(G)}$. The map L is a bijection between the open cone containing G and the set $(0, \infty)^{E(G)}$. By projectivizing, we obtain a bijection between the open cell of G in \mathcal{X}_r and the interior of the simplex $\mathcal{P}[0, \infty)^{E(G)}$.

The *closed cone* in $\hat{\mathcal{X}}_r$ containing G is the set of elements represented by marked graphs G' for which there exists a quotient map $p: G \rightarrow G'$ that preserves marking and that collapses to a point each component of a subforest $F \subset G$; the *closed cell* in \mathcal{X}_r containing G is similarly defined. The map p induces a well-defined bijection $E(G) - E(F) \leftrightarrow E(G')$, and therefore there is a well-defined vector $L_{G'} \in [0, \infty)^{E(G)}$ which assigns zero to edges in F and assigns the length of the corresponding edge of G' to each edge in $G - F$. Again L is a bijection between the closed cone containing G and a certain subset of $[0, \infty)^{E(G)}$, namely those vectors which assign coordinate zero to a subset of $E(G)$ that forms a subforest of G . The closed cone of G is thus identified with a subset of $[0, \infty)^{E(G)}$ which is the union of the interior and certain of the faces. By projectivizing, we obtain a bijection between the closed cell of G in \mathcal{X}_r and a union of faces of the simplex $\mathcal{P}[0, \infty)^{E(G)}$. Note that if $F \subset F'$ are faces of the simplex and F is a face of the closed cone of G then F' is also a face of the closed cone of G .

The *cellular topology* on $\widehat{\mathcal{X}}_r$ is the weakest topology such that for each marked graph G , the closed cone containing G is a closed subset of $\widehat{\mathcal{X}}_r$, and the map L from this cone to $[0, \infty)^{E(G)}$ is a homeomorphism onto its image. The cellular topology on \mathcal{X}_r is the quotient by projectivization of the cellular topology on $\widehat{\mathcal{X}}_r$.

We need some facts about the cellular topology. First, the closed cone (simplex) of G' is a face of the closed cone (simplex) of G if and only if there is a quotient map $G \mapsto G'$ that preserves marking and that collapses to a point each component of some subforest of G . Second, the closed cells (simplices) are locally finite: for any G' there are only finitely many closed cells (simplices) that contain G' , and hence by the weak topology there is a neighborhood of G' that intersects only those finitely many cells (simplices) that contain G' .

Given a marked graph G define its *length* $\ell(G)$ to be the sum of the lengths of its edges. This gives a well-defined function on $\widehat{\mathcal{X}}_r$ which, on each closed cell, is just the sum of the coordinates. It follows that, in terms of the cellular topology:

Proposition 2.1. *The map $\ell: \widehat{\mathcal{X}}_r \rightarrow (0, \infty)$ is continuous.* \diamond

The length or axes topology [CM87]. Let \mathcal{C} be the set of nontrivial conjugacy classes in F_r , give $\mathbf{R}^{\mathcal{C}}$ the product topology and let $PR^{\mathcal{C}}$ be the associated projective space. For each marked graph G , each $c \in \mathcal{C}$ corresponds bijectively, via the marking ρ_G , with an immersion $S^1 \rightarrow G$ also denoted c , which is well-defined up to a homeomorphism of the domain, and whose length $\text{Len}_G(c)$ is well-defined. The map $G \mapsto \text{Len}_G$ defines an injection $\widehat{\mathcal{X}}_r \hookrightarrow \mathbf{R}^{\mathcal{C}}$ and an induced injection $\mathcal{X}_r \hookrightarrow PR^{\mathcal{C}}$. These injections may also be defined in terms of translations lengths of free, simplicial F_r -trees representing elements of \mathcal{X}_r . We give \mathcal{X}_r and $\widehat{\mathcal{X}}_r$ the topologies that make these injections homeomorphisms onto their images; we shall occasionally refer to these as the *length topologies* on \mathcal{X}_r and $\widehat{\mathcal{X}}_r$. It follows that the natural action of \mathbf{R}_+ on $\widehat{\mathcal{X}}_r$, where $t \cdot G$ is obtained from G by multiplying the metric by the constant factor t , is an action by homeomorphisms with quotient space \mathcal{X}_r .

The Gromov topology [Pau89]. This topology on $\widehat{\mathcal{X}}_r$ is defined by the following basis. Given an F_r -tree $T \in \widehat{\mathcal{X}}_r$, $\epsilon > 0$, a finite subset $K \subset T$, and a finite subset $P \subset F_r$, define the basis element $V(T; \epsilon, K, P)$ to be the set of F_r -trees $T' \in \widehat{\mathcal{X}}_r$ for which there exists a finite subset $K' \subset T'$ and a relation $R \subset K \times K'$ such that R projects onto K and onto K' , and the following hold:

R is an ϵ -almost isometry: For all $x, y \in K$ and $x', y' \in K'$, if $(x, x'), (y, y') \in R$ then $|d(x, y) - d(x', y')| < \epsilon$.

R is P -equivariant: For all $x \in K$, $x' \in K'$, and $\alpha \in P$, if $\alpha \cdot x \in K$ and $(x, x') \in R$ then $\alpha \cdot x' \in K'$ and $(\alpha \cdot x, \alpha \cdot x') \in R$.

In [Pau89] the Gromov topology is defined by forming a basis as above but using compact subsets of T and T' instead of finite subsets; as shown in Proposition 4.1 of [Pau89] the definition given above is a basis for the same topology.

The quotient topology on \mathcal{X}_r is also called the Gromov topology.

The following theorem is a combination of results of [CV86] and [Pau89]:

Theorem 2.2. *The cellular topology, the length topology, and the Gromov topology are identical on $\widehat{\mathcal{X}}_r$ and on \mathcal{X}_r .* \diamond

Compactification and boundary of outer space. The closure of \mathcal{X}_r in PR^C is denoted $\overline{\mathcal{X}}_r$. Its boundary is $\partial\mathcal{X}_r = \overline{\mathcal{X}}_r - \mathcal{X}_r$. Both $\overline{\mathcal{X}}_r$ and $\partial\mathcal{X}_r$ are compact [CM87]. Let $\widehat{\mathcal{X}}_r$ be the inverse image of $\overline{\mathcal{X}}_r$ under the projection $\mathbf{R}^C - \{O\} \mapsto PR^C$, equivalently, the closure of $\widehat{\mathcal{X}}_r$ in $\mathbf{R}^C - \{O\}$. Let $\partial\widehat{\mathcal{X}}_r = \widehat{\mathcal{X}}_r - \mathcal{X}_r$, which is just the inverse image under projection of $\partial\mathcal{X}_r$.

We next relate compactified outer space to actions on trees.

Theorem 2.3. (*Culler–Morgan [CM87]*) *For each $\ell \in \widehat{\mathcal{X}}_r$ there exists an F_r -tree T such that $\ell = \ell_T$, and T depends uniquely on ℓ up to equivariant isometry.*

Extending our notation, when $\ell = \ell_T$ as in this theorem, we use $[T]$ to denote the corresponding point in $\widehat{\mathcal{X}}_r$ or $\overline{\mathcal{X}}_r$ as the context makes clear.

Lemma 2.4. *Given $\ell \in \widehat{\mathcal{X}}_r$ and given T with $\ell_T = \ell$, the following are equivalent:*

- (1) $\ell \in \widehat{\mathcal{X}}_r$.
- (2) $\ell(c)$ is bounded away from zero for $c \in \mathcal{C}$.
- (3) The action of F_r on T is free and properly discontinuous.

Proof. It is obvious that the first item implies the second. Since the action is by isometries the second item implies the third. Finally, the third item implies that the quotient map from T to $G := T/F_r$ is a covering space. Since T is an F_r -tree, G is a finite graph and so T is simplicial. \diamond

2.2 Paths, Circuits and Edge Paths

Consider a marked graph G and the universal covering projection $pr : \widetilde{G} \rightarrow G$. A proper map $\tilde{\sigma} : J \rightarrow \widetilde{G}$ with domain a (possibly infinite) interval J will be called a *path in \widetilde{G}* if it is an embedding or if J is finite and the image is a single point; in the latter case we say that $\tilde{\sigma}$ is a *trivial path*. If J is finite, then every map $\tilde{\sigma} : J \rightarrow \widetilde{G}$ is homotopic rel endpoints to a unique (possibly trivial) path $[\tilde{\sigma}]$; we say that $[\tilde{\sigma}]$ is *obtained from $\tilde{\sigma}$ by tightening*.

We will not distinguish between paths in \widetilde{G} that differ only by an orientation preserving change of parametrization. Thus we are interested in the oriented image of $\tilde{\sigma}$ and not $\tilde{\sigma}$ itself. If the domain of $\tilde{\sigma}$ is finite, then the image of $\tilde{\sigma}$ has a natural decomposition as a concatenation $\tilde{E}_1 \cdot \tilde{E}_2 \cdot \dots \cdot \tilde{E}_k$, where \tilde{E}_i for $1 < i < k$ is an edge of \widetilde{G} , \tilde{E}_1 is the terminal segment of some edge, and \tilde{E}_k is the initial segment of some edge. If the endpoints of the image of $\tilde{\sigma}$ are vertices, then \tilde{E}_1 and \tilde{E}_k are full edges. The sequence $\tilde{E}_1 \cdot \tilde{E}_2 \cdot \dots \cdot \tilde{E}_k$ is called the *edge path associated to $\tilde{\sigma}$* , and the number k is called the *combinatorial length* of this edge path. When the interval of domain is half-infinite then an edge path has the form $\tilde{E}_1 \cdot \tilde{E}_2 \cdot \dots$ or $\dots \cdot \tilde{E}_{-2} \cdot \tilde{E}_{-1}$, and in the bi-infinite case an edge path has the form $\dots \cdot \tilde{E}_{-1} \cdot \tilde{E}_0 \cdot \tilde{E}_1 \cdot \tilde{E}_2 \cdot \dots$.

A *path in G* is the composition of the projection map pr with a path in \widetilde{G} . Thus a map $\sigma : J \rightarrow G$ with domain a (possibly infinite) interval will be called a *path* if it is an immersion or if J is finite and the image is a single point; paths of the latter type are said to be trivial. If J is finite, then every map $\sigma : J \rightarrow G$ is homotopic rel endpoints to a unique (possibly trivial) path $[\sigma]$; we say that $[\sigma]$ is *obtained from σ by tightening*. For any lift $\tilde{\sigma} : J \rightarrow \Gamma$

of σ , $[\sigma] = pr[\tilde{\sigma}]$. We do not distinguish between paths in G that differ by an orientation preserving change of parametrization. The *edge path associated to σ* is the projected image of the edge path associated to a lift $\tilde{\sigma}$. Thus the edge path associated to a path with finite domain has the form $E_1 \cdot E_2 \cdot \dots \cdot E_{k-1} \cdot E_k$ where E_i , $1 < i < k$, is an edge of G , E_1 is the terminal segment of an edge, E_k is the initial segment of an edge; the *combinatorial length* again equals k .

We reserve the word *circuit* for an immersion $\sigma : S^1 \rightarrow G$. Any homotopically non-trivial map $\sigma : S^1 \rightarrow G$ is homotopic to a unique circuit $[\sigma]$. As was the case with paths, we do not distinguish between circuits that differ only by an orientation preserving change in parametrization and we identify a circuit σ with a *cyclically ordered edge path* $E_1 E_2 \dots E_k$, and again the *combinatorial length* equals k .

If $h : H \rightarrow G$ is a homotopy equivalence of marked graphs and if $\tilde{h} : \tilde{H} \rightarrow \tilde{G}$ is a lift then for each path $\tilde{\sigma}$ in \tilde{H} we denote $[\tilde{h}(\tilde{\sigma})]$ by $\tilde{h}_\#(\tilde{\sigma})$, and for each path or circuit σ in H we denote $[h(\sigma)]$ by $h_\#(\sigma)$. Given a self homotopy equivalence $g : G \rightarrow G$ of a marked graph and a path or circuit σ in G , a decomposition $\sigma = \dots \cdot \sigma_i \cdot \sigma_{i+1} \cdot \dots$ is called a *splitting* if $g_\#^k(\sigma) = \dots \cdot g_\#^k(\sigma_i) \cdot g_\#^k(\sigma_{i+1}) \cdot \dots$ for all $k \geq 0$; in other words, to tighten $g^k(\sigma)$ one need only tighten the $g^k(\sigma_i)$.

Throughout this paper we will identify paths and circuits with their associated edge paths.

A path or circuit *crosses* or *contains* an edge if that edge occurs in the associated edge path. For any path σ in G define $\bar{\sigma}$ to be ‘ σ with its orientation reversed’. For notational simplicity, we sometimes refer to the inverse of $\tilde{\sigma}$ by $\tilde{\sigma}^{-1}$.

Bounded cancellation. The bounded cancellation lemma was introduced in [Coo87] and generalized to the form we need in [BFH97]. Let $\text{Lip}(f)$ be the lipschitz constant of a map f .

Lemma 2.5. *[Bounded Cancellation, Lemma 3.1 of [BFH97]] Let T_0 be a free, simplicial F_r -tree, let T be an F_r -tree representing a point in $\overline{\mathcal{X}}_r$, and let $f : T_0 \rightarrow T$ be an F_r -equivariant map. There exists a constant $\text{BCC}(f) \leq \text{Lip}(f) \cdot \text{Len}(T_0/F_r)$ such that for each geodesic segment $\alpha \subset T_0$, the image $f(\alpha)$ is contained in the $\text{BCC}(f)$ neighborhood of the straightened image $f_\#(\alpha)$.*

2.3 Folds

In this subsection we define morphisms, edge isometries, and fold paths. We also define fold maps and Stallings fold sequences, and we describe relations between these objects and fold paths.

Turns and gates. In any graph a *turn* is an unordered pair of directions with a common base point. The turn is *nondegenerate* if it is defined by two distinct directions and is *degenerate* otherwise.

Consider two graphs T, T' and a map $H : T \rightarrow T'$ which is locally injective on each edge. For each $x \in T$ there is an induced map $D_x H$ from the set of directions of T at x to the set of directions of T' at $H(x)$: given a direction d at x represented by an oriented path α with

initial endpoint x , $D_x H(d)$ is defined to be the direction at $H(x)$ represented by $H(\alpha)$; the map $D_x H$ is well-defined, independent of the choice of α representing d . We use DH either as an abbreviation for $D_x H$ when x is understood, or as the union of all of the maps $D_x H$ over all $x \in T$.

Given an F_r -equivariant map of F_r -trees $H: T \rightarrow T'$, we define the *gates* at a point $x \in T$ with respect to H to be the equivalence classes of directions at x where d, d' are equivalent if $D_x H(d) = D_x H(d')$. In applications where the base point x of a direction is a vertex we can identify directions with oriented edges E in the link of x , and if in addition H takes vertices to vertices then $D_x H(E)$ is the first edge in the edge path associated to $H(E)$. We have here extended the terminology of a “gate” which in the literature has so far been applied only to train track maps; see Section 2.4.

Piecewise isometries and edge isometries. An F_r -equivariant map $f: T \rightarrow T'$ from a free, simplicial F_r -tree T to an F_r tree T' is a *piecewise isometry*² if each edge of T may be subdivided into finitely many segments on each of which f is an isometry. We usually use the even stronger property that f restricted to each edge of T is an isometry, to which we refer by saying that f is an *edge isometry* or that f is *edge isometric*. For a piecewise isometry to be an edge isometry is equivalent to it being locally injective on each open edge. When T' is also free and simplicial then we also define a piecewise isometry or edge isometry of marked graphs $G = T/F_r \mapsto T'/F_r = G'$, meaning that the lifted map $T \rightarrow T'$ is a piecewise isometry or edge isometry. Equivalently, an edge isometry between marked graphs is any homotopy equivalence that preserves marking and restricts to a local isometry on each open edge.

Fold paths. A *fold path* in outer space is a continuous, proper injection $\gamma: I \rightarrow \mathcal{X}_r$, where $I \subset \mathbf{R}$ is connected, so that there exists a continuous lift $\hat{\gamma}: I \rightarrow \hat{\mathcal{X}}_r$, a marked graph G_r representing $\hat{\gamma}(r)$ for each $r \in I$, and maps $g_{sr}: G_r \rightarrow G_s$ for each $r \leq s$ in I , satisfying the following:

- (1) **Train track property:** Each g_{sr} is an edge isometry.
- (2) **Semiflow property:** $g_{tr} = g_{ts} \circ g_{sr}$ for each $r \leq s \leq t$ in I , and $g_{ss}: G_s \rightarrow G_s$ is the identity for each $s \in I$.

A fold path γ whose domain I is noncompact on both ends is called a *fold line*.

Fold paths are investigated in detail in Section 7. For now we turn to well known constructions that lead to simple examples of fold paths.

Fold maps. Consider a marked graph G and two distinct oriented edges e, e' of G that have the same initial vertex v . Parameterize e, e' by arc length from their initial vertex, using the notation $e(\cdot), e'(\cdot)$. Choose a number S such that $e([0, S]) \cap e'([0, S]) = \{v\}$. For $s \in [0, S]$ define G_s to be the graph obtained from G by identifying $e(t)$ to $e'(t)$ for each $t \in [0, s]$. The quotient map $G \mapsto G_s$ is a homotopy equivalence, because of the fact that $e(s) \neq e'(s)$. Push forward the metric on edges of G by the quotient map $G \mapsto G_s$, defining a metric on edges

²A “piecewise isometry” is the same thing as a “morphism”, as commonly defined in the literature, but restricted so that the domain is free and simplicial.

of G_s which extends naturally to a geodesic metric on all of G_s . Postcompose the marking of G by the homotopy equivalence $G \mapsto G_s$, defining a marking of G_s . It follows that the quotient map $G \mapsto G_s$ preserves the marking and is an edge isometry; when $s = 0$ this is just the identity map on G . When $s > 0$ we say that G_s is obtained from G by a *length s fold of e and e'* , and the quotient map $G \mapsto G_s$ is called a *fold map* from G to G_s .

Consider the map $\hat{\gamma}: [0, S] \rightarrow \hat{\mathcal{X}}_r$ defined by $\hat{\gamma}(s) = G_s$. First, $\hat{\gamma}$ is a continuous path in $\hat{\mathcal{X}}_r$ from G to G_S , and so it descends to a continuous path $\gamma: [0, S] \rightarrow \mathcal{X}_r$ which, its domain being compact, is proper. Second, for $0 \leq s \leq t \leq S$ the induced map $f_{ts}: G_s \mapsto G_t$ is a fold map. Third, the family of maps f_{ts} satisfies the train track property and the semiflow property, by construction. It follows that γ is a fold path. Either of the paths $\gamma: [0, S] \rightarrow \mathcal{X}_r$ or $\hat{\gamma}: [0, S] \rightarrow \hat{\mathcal{X}}_r$, or any of their orientation preserving reparameterizations, are called *interpolating paths* or just *interpolations* of the fold map $G \mapsto G_S$.

Stallings fold sequences. Given an edge isometry $h: G \rightarrow H$ between marked graphs, a *Stallings fold sequence* for h is a finite sequence of edge isometric fold maps

$$G = G_0 \xrightarrow{f_{1,0}} G_1 \xrightarrow{f_{2,1}} \cdots \xrightarrow{f_{J,J-1}} G_J = H \quad (2.2)$$

$\underbrace{\hspace{10em}}_h$

whose composition equals h . If we define f_{ji} for integer values $0 \leq i \leq j \leq J$ to be the composition $f_{j,j-1} \circ \cdots \circ f_{i+1,i}: G_i \rightarrow G_j$, then the maps f_{ji} are edge isometries satisfying the semiflow property $f_{ki} = f_{kj} \circ f_{ji}$ for integer values $0 \leq i \leq j \leq k \leq J$. For $j = 1, \dots, J$ let $[j-1, j] \mapsto \hat{\mathcal{X}}_r$ be an interpolating path for the fold map $f_{j,j-1}: G_{j-1} \rightarrow G_j$, and extend the map $f_{j,j-1}$ to a family of edge isometries $f_{ts}: G_s \rightarrow G_t$, $j-1 \leq s \leq t \leq j$ that satisfies the semiflow property. Concatenating the paths $[j-1, j] \mapsto \hat{\mathcal{X}}_r$ for $j = 1, \dots, J$ defines a continuous path $\hat{\gamma}: [0, J] \mapsto \hat{\mathcal{X}}_r$, and the entire family of edge isometries f_{ts} defined for $0 \leq s \leq t \leq J$ satisfies the semiflow property. Thus $\hat{\gamma}$ descends to a continuous path $\gamma: [0, J] \rightarrow \mathcal{X}_r$ which is a fold path, called the *Stallings fold path* of the Stallings fold sequence (2.2).

To complete the construction of Stallings fold paths, we review Stallings original topological construction of fold sequences [Sta83], and adapt this construction to our present metric setting where fold maps are required to be edge isometric. At the same time we shall use this construction to bound the length of the fold sequence.

In a Stallings fold sequence (2.2), the fold $f_{j+1,j}: G_j \rightarrow G_{j+1}$ is said to be *complete* if it satisfies the following: assuming $f_{j+1,j}$ to be a length S fold of oriented edges e, e' of G_j , then S is equal to the maximum value of s such that restrictions of the map $f_{J,j}: G_j \rightarrow G_J = H$ to $e([0, s])$ and to $e'([0, s])$ are the same path in H . In other words, $f_{j+1,j}$ is complete if the maximal possible initial segments of e, e' are folded immediately.

Define the *combinatorial length* of an edge isometry $h: G \rightarrow H$ as follows. While at first regarding G and H with their natural vertex sets consisting solely of vertices of valence ≥ 3 , subdivide H at the images of the vertices of G , which adds at most $2r - 2$ new valence 2 vertices to H and at most $2r - 2$ new edges, bringing the number of edges in H up to at most $5r - 5$. For each edge e of G we may regard $h(e)$ as an edge path in the subdivided graph H , and we define the *combinatorial length* of h to be the sum of the combinatorial lengths of $h(e)$ over all edges e of G . Equivalently, having subdivided H , if we then subdivide G at the

inverse image of the vertex set of H , and if we refer to the edges of this subdivision as the *edgelets* of G with respect to h , then the combinatorial length of h is the number of edgelets of G with respect to h .

Lemma 2.6. *Each edge isometry $h: G \rightarrow H$ which is not an isometry has a sequence of complete Stallings folds. The length of any such fold sequence is bounded above by the combinatorial length of h .*

Proof. Subdivide H at the images of the vertices of G .

Since the map h is a local injection on each edge and the graphs G, H have no valence 1 vertices, if h were locally injective at each vertex then there would be an extension of h to a covering map $h': G' \rightarrow H$ where the graph G' deformation retracts to G (see [Sta83]). Since h is a homotopy equivalence it would follow that $h = h'$ is a homeomorphism and therefore h is an isometry, contrary to the hypothesis.

There exist therefore two oriented edges e, e' of G with the same initial vertex v , and there exists a positive number $R \leq \min\{\text{Len}(e), \text{Len}(e')\}$, such that h maps $e([0, R])$ and $e'([0, R])$ to the same oriented path in H . The maximal such R has the property that $h(e(R)) = h(e'(R))$ is a vertex of H and $e([0, R]) \cap e'([0, R]) = \{v\}$. Define $G \xrightarrow{h'} G'$ to be the length R fold of e and e' , and so h factors as $G \xrightarrow{h'} G' \xrightarrow{h''} H$. Noting that each of $e([0, R])$ and $e'([0, R])$ is a union of edgelets of G with respect to h , identifying $e([0, R])$ with $e'([0, R])$ has the effect of identifying their edgelets in pairs, and so the number of edgelets of G' with respect to the edge isometry h'' — the combinatorial length of h'' — is strictly less than the number of edgelets of G with respect to h — the combinatorial length of h . The existence of a sequence of complete Stallings folds now follows by induction on combinatorial length.

The same argument shows that for *any* sequence of complete Stallings folds (2.2), the combinatorial length of $f_{J,j}: G_j \rightarrow G_J = H$ is a strictly decreasing function of j , proving the desired bound. \diamond

2.4 Train Track Maps

A topological representative $g: G \rightarrow G$ of an outer automorphism ϕ is an *irreducible train track map* if $g^k|_E$ is an immersion for each edge E and each $k > 0$ and if for every pair of edges E and E' there exists $l > 0$ such that $g^l(E)$ crosses E' . Theorem 1.7 of [BH92] states that every irreducible $\phi \in \text{Out}(F_r)$ is represented by an irreducible train track map. Moreover, there exists $\lambda > 1$, called the *expansion factor* for ϕ , there exists a metric on G , and there exists a homotopy of g so that, still using g for the homotoped map, for all edges E the restriction $g|_E$ uniformly expands length by the factor λ . We will always assume that the graph G is equipped with such a metric, and we occasionally emphasize this point by referring to g as an *affine train track map*.

A nondegenerate turn in G is *illegal* with respect to $g: G \rightarrow G$ if its image under some iterate of Dg is degenerate, and otherwise the turn is *legal*. A path $\sigma = E_1 \cdot E_2 \cdot \dots \cdot E_{k-1} \cdot E_k$ contains or takes each of the turns $\{\bar{E}_i, E_{i+1}\}$ for $0 \leq i \leq k-1$. A path $\sigma \subset G$ is *legal* if it contains only legal turns. It is immediate from the definitions that Dg maps legal turns to legal turns and that g maps legal paths to legal paths. Note that g expands the length of any legal subpath by the factor λ .

We define the *gates* of a train track map $g: G \rightarrow G$ as follows. The map g induces a map Dg on the set of directions at vertices of g . The *gates* at a vertex x with respect to H are the equivalence classes of directions at x where d, d' are equivalent if there exists $i \geq 1$ such that $(Dg)^i(d) = (Dg)^i(d')$. It follows that two directions are in the same gate if and only if the turn they make is illegal. Note the distinction between gates of a train track map and gates of an F_r -equivariant map of F_r -trees which is defined in Section 2.3.

Local Whitehead Graphs. Consider an irreducible train track representative $g: G \rightarrow G$ of an irreducible $\phi \in \text{Out}(F_r)$. For each point $x \in G$, the *local Whitehead graph at x* , denoted $W(x)$ or more explicitly $W(x; G)$, is the graph whose vertices are the directions d_i at x and with an edge connecting d_i to d_j if there exists an edge E and $k > 0$ such that the path $g^k(E)$ takes the turn (d_i, d_j) . Since ϕ is irreducible, $W(x)$ is connected (see section 2 of [BFH97]). For each $x \in G$, the map $D_x g$ extends to a simplicial map $W(x) \mapsto W(g(x))$ also denoted $D_x g$, because each turn at x taken by some $g^k(E)$ maps to a turn at $g(x)$ taken by $g^{k+1}(E)$.

If $x \in \text{Per}(g)$, then the *local stable Whitehead graph at x* , denoted $SW(x)$ or $SW(x; G)$, is the subgraph of $W(x)$ obtained by restricting to the periodic directions at x and the edges between them. Since some iterate g^k induces a continuous map $D_x g^k$ from $W(x)$ onto $SW(x)$, it follows that $SW(x)$ is connected. There are analogous definitions of $W(\tilde{x}) = W(\tilde{x}; \tilde{G})$ and $SW(\tilde{x}) = SW(\tilde{x}; \tilde{G})$ for any vertex \tilde{x} of \tilde{G} ; these can also be defined as the lifts of $W(x)$ and $SW(x)$. It follows that $W(\tilde{x})$ and $SW(\tilde{x})$ are connected.

Nielsen Paths. A path σ is a *periodic Nielsen path* if $g_\#^p(\sigma) = \sigma$ for some minimal $p \geq 1$ called the *period* of σ . A *Nielsen path* is a periodic Nielsen path with period one. A periodic Nielsen path is *indivisible* if it can not be written as a non-trivial concatenation of periodic Nielsen subpaths. A path σ is a *periodic pre-Nielsen path* if it is not a periodic Nielsen path but some $g_\#^k(\sigma)$ is a periodic Nielsen path.

We record information about Nielsen paths in the following lemma.

Lemma 2.7. *Assume that $g: G \rightarrow G$ is an irreducible train track map.*

- (1) *There are only finitely many indivisible periodic or pre-periodic Nielsen paths for g .*
- (2) *An indivisible periodic or pre-periodic Nielsen path σ decomposes into subpaths $\sigma = \alpha\bar{\beta}$ where α and β are legal and the turn $(\bar{\alpha}, \beta)$ is illegal. Moreover, if α_n and β_n are the terminal segments of α and β that are identified by g^n , then the $\alpha_n\bar{\beta}_n$'s are an increasing sequence of subpaths of σ whose union is the interior of σ .*
- (3) *For every path or circuit σ there exists $K(\sigma) \geq 0$ such that for all $k \geq K(\sigma)$, there is a splitting of $g_\#^k(\sigma)$ into subpaths that are either legal or are indivisible periodic Nielsen paths.*
- (4) *There exists $K^*, L^* > 0$ so that if σ is a path that decomposes as the concatenation of two legal subpaths each of length at least L^* then $K(\sigma) \leq K^*$.*
- (5) *If σ is a bi-infinite path with exactly one illegal turn then one of the following holds for all $k \geq K^*$.*

- $g_{\#}^k(\sigma)$ is legal.
- $g_{\#}^k(\sigma)$ splits as $R_k \cdot \tau_k \cdot R'_k$ where R_k and R'_k are legal rays and τ_k is an indivisible periodic Nielsen path.

Proof. The first three items follow from Lemma 4.2.5 and Lemma 4.2.6 of [BFH00]. For the fourth item, let

$$L^* = \text{Len}(G) + \min\{\text{Len}(e) \mid e \text{ is an edge of } G\}$$

If σ is legal then $K(\sigma) = 1$ so we may assume that σ has exactly one illegal turn. Decompose σ as $\sigma = \sigma_1 \bar{\sigma}_2$ where σ_1 and σ_2 are legal and let $\sigma_{1,n}$ and $\sigma_{2,n}$ be the terminal segments of σ_1 and σ_2 that are identified by g^n . If there does not exist $n > 0$ such that $\sigma_{i,k} = \sigma_{i,n}$ for all $k \geq n$ then the closure of $\cup_n \sigma_{1,n} \bar{\sigma}_{2,n}$ is an indivisible Nielsen or pre-Nielsen path by Lemma 4.2.5 of [BFH00] and σ splits into an initial legal path, followed by $\cup_n \sigma_{1,n} \bar{\sigma}_{2,n}$ followed by a terminal legal path. By item (1), the pre-Nielsen paths become periodic Nielsen paths after a uniformly bounded number of iterations, and so $K(\sigma)$ is uniformly bounded. We may therefore assume that there is such an n . By Lemma 2.5 (see also Lemma 3.9 of [BH92]) the length of $\sigma_{i,n}$ is bounded above by $\text{Len}(G)$. Removing an initial segment of σ_i , $i = 1, 2$, that is disjoint from $\sigma_{i,n}$ does not change $K(\sigma)$. Since $\text{Len}(\sigma_i) \geq L^*$ we may therefore remove such initial segments so that σ has endpoints at vertices and has length at most $2\text{Len}(G) + 2\max\{\text{Len}(e) \mid e \text{ is an edge of } G\}$. There are only finitely many such σ and we take K^* to be the maximum of their $K(\sigma)$'s. This completes the proof of (4).

Lemma 4.2.2 of [BFH00] implies that a bi-infinite line with exactly one illegal turn splits as $R \cdot \tau \cdot R'$ where R and R' are legal rays and τ decomposes as the concatenation of two legal subpaths each of length at least L^* . The fifth item therefore follows from the third and fourth items and the fact that $g_{\#}$ maps legal paths to legal paths and indivisible periodic Nielsen paths to indivisible periodic Nielsen paths. \diamond

Assumption: We assume from now on that for any irreducible train track map $g: G \rightarrow G$, the endpoints of indivisible periodic Nielsen paths are vertices. This can always be arranged by a finite subdivision.

Principal vertices, principal lifts and rotationless train track maps.

We now recall from [FH04] the concepts of principal automorphisms, principal lifts of train track representatives, and rotationless outer automorphisms. While we apply these concepts here only to fully irreducible outer automorphisms, more general definitions are given in [FH04] which apply to arbitrary outer automorphisms and their relative train track representatives.

Let $\phi \in \text{Out}(F_r)$ be fully irreducible, and consider $\Phi \in \text{Aut}(F_r)$ representing ϕ and its boundary extension $\hat{\Phi}: \partial F_r \rightarrow \partial F_r$. Denote the fixed point set of $\hat{\Phi}$ by $\text{Fix}(\hat{\Phi})$, and denote the subset of nonrepelling fixed points of $\text{Fix}(\hat{\Phi})$ by $\text{Fix}_N(\hat{\Phi})$. We say that Φ is a *principal automorphism* if $\text{Fix}_N(\hat{\Phi})$ contains at least three points. The set of principal automorphisms representing ϕ is denoted $\text{PA}(\phi)$. Two elements $\Phi, \Phi' \in \text{PA}(\phi)$ are *equivalent*³ if there exists $\gamma \in F_n$ such that $\Phi' = i_{\gamma} \circ \Phi \circ i_{\gamma}^{-1}$.

If the fixed subgroup $\text{Fix}(\Phi)$ is non-trivial, then $\text{Fix}(\Phi) = \langle \gamma \rangle$ is infinite cyclic. This happens if and only if ϕ is geometric, in which case ϕ is represented by a pseudo-Anosov

³This equivalence relation is called “isogredience” in [Nie86] and [LL00a].

homeomorphism of a once punctured surface whose boundary is represented by γ (see Theorem 4.1 of [BH92]).

The following lemma is contained in Lemma 2.3 and Lemma 2.4 of [BFH04] and in Proposition 1.1 of [GJLL98].

Lemma 2.8. *Assume that $\Phi \in \text{PA}(\phi)$.*

- (1) *If $\text{Fix}(\Phi)$ is trivial then $\text{Fix}_N(\hat{\Phi})$ is a finite set of attractors.*
- (2) *If $\text{Fix}(\Phi) = \langle \gamma \rangle$ is infinite cyclic, then $\text{Fix}_N(\hat{\Phi})$ is the union of t_γ^\pm with a finite set of t_γ -orbits of attractors.*
- (3) *If $P \in \text{Fix}_N(\hat{\Phi})$ is an attractor then it is not t_γ^+ or t_γ^- for any nontrivial covering translation t_γ .*

Corollary 2.9. *If $\Phi_1 \neq \Phi_2 \in \text{PA}(\phi)$ then $\text{Fix}(\hat{\Phi}_1)$ and $\text{Fix}(\hat{\Phi}_2)$ are disjoint.*

Proof. Suppose that $P \in \text{Fix}(\hat{\Phi}_1) \cap \text{Fix}(\hat{\Phi}_2)$. Observe that $\Phi_1^{-1}\Phi_2$ is a covering translation t_γ , for a nontrivial $\gamma \in F_r$. It follows that $t_\gamma(P) = P$ and hence that $\gamma \in \text{Fix}(\Phi_1) \cap \text{Fix}(\Phi_2)$. In this case, ϕ is geometric and the corollary is well known; see for example Lemma 4.2 of [HT85]. \diamond

Let $\phi \in \text{Out}(F_r)$ be fully irreducible. We say that ϕ is *rotationless* if for each $k > 0$ and each principal automorphism Φ_k representing ϕ^k there is a principal automorphism Φ representing ϕ such that $\Phi_k = \Phi^k$ and such that $\text{Fix}_N(\hat{\Phi}) = \text{Fix}_N(\hat{\Phi}_k)$. In other words, the set of principal automorphisms and their nonrepelling fixed points in ∂F_r is stable under iteration.

Let $g: G \rightarrow G$ be an irreducible train track representative of ϕ . A lift $\check{g}: \tilde{G} \rightarrow \tilde{G}$ is a *principal lift* if the corresponding automorphism is a principal automorphism, and so two principal lifts \check{g}, \check{g}' are equivalent if and only if $\check{g}' = t_\gamma \circ \check{g} \circ t_\gamma^{-1}$ for some $\gamma \in F_r$. Consider a periodic vertex $v \in G$. Since v has at least one legal turn, v has at least two periodic directions. We say that v is a *principal vertex* if there are at least three periodic directions at v or if v is the endpoint of an indivisible periodic Nielsen path. The set of principal vertices is non-empty by Lemma 3.18 of [FH04]. If every principal vertex $v \in G$ is fixed and if every periodic direction at a principal vertex is fixed then we say that $g: G \rightarrow G$ is a *rotationless train track map*.

These definitions are connected by Proposition 3.24 of [FH04], which we recall.

Proposition 2.10. *Assume that $g: G \rightarrow G$ is an irreducible train track map for a fully irreducible $\phi \in \text{Out}(F_r)$. Then ϕ is rotationless if and only if $g: G \rightarrow G$ is rotationless.*

Remark. It is clear that g^p , and hence ϕ^p , is rotationless for some $p \geq 1$. Some of our results are stated for rotationless ϕ and $g: G \rightarrow G$; to apply these in general one replaces ϕ and g by suitable iterates.

A pair of points $x, y \in \text{Fix}(g)$ are *Nielsen equivalent* if they are the endpoints of a Nielsen path in G . The elements of a Nielsen equivalence class are usually referred to simply as a *Nielsen class*. It is well known, and easy to check, that x and y are Nielsen equivalent if and

only if there is a lift $\check{g} : \tilde{G} \rightarrow \tilde{G}$ and points $\tilde{x}, \tilde{y} \in \text{Fix}(\check{g})$ that project to x and y respectively. In particular, for any lift \check{g} , the image of $\text{Fix}(\check{g})$ in G is either empty or is an entire Nielsen class. Moreover, if $\text{Fix}(\check{g}_1)$ and $\text{Fix}(\check{g}_2)$ are non-empty, then \check{g}_1 and \check{g}_2 are equivalent, as defined above in the context of principal lifts, if and only if $\text{Fix}(\check{g}_1)$ and $\text{Fix}(\check{g}_2)$ project to the same Nielsen class in $\text{Fix}(g)$. It is useful to note that if $\text{Fix}(\check{g}_1) \cap \text{Fix}(\check{g}_2) \neq \emptyset$ then $\check{g}_1 = \check{g}_2$. This follows from the fact that \check{g}_1 and \check{g}_2 differ by a covering translation and the fact that every non-trivial covering translation is fixed point free.

It will be convenient to refer to lifts $\tilde{v} \in \tilde{G}$ of principal vertices $v \in G$ as *principal vertices of \tilde{G}* .

Lemma 2.11. *Assume that $g : G \rightarrow G$ is rotationless. Then $\text{PA}(\phi) \neq \emptyset$. Moreover, \check{g} is a principal lift if and only if $\text{Fix}(\check{g}) \neq \emptyset$ and some, and hence every, element of $\text{Fix}(\check{g})$ is a principal vertex.*

Proof. The second statement follows from Corollary 3.14 and Corollary 3.22 of [FH04]. The first follows from the second and Lemma 3.18 of [FH04]; see also Proposition I.5 of [LL00b]. \diamond

We also recall Lemma 3.23 of [FH04].

Lemma 2.12. *If g is rotationless then every periodic Nielsen path has period one.* \diamond

2.5 The attracting tree T_+

Given a fully irreducible $\phi \in \text{Out}(F_r)$ there exists a unique point $T_+ = T_+^\phi \in \partial \mathcal{X}_r$ which is the attracting point for every forward orbit of ϕ in \mathcal{X}_r ; this is proved in [BFH00]. We review here the construction and properties of the F_r -tree T_+ . Although we shall not need this fact here, we mention the result of [LL03] that the action of ϕ on \bar{X}_r has source-sink dynamics with sink T_+^ϕ and source $T_-^\phi = T_+^{\phi^{-1}}$.

Let $g : G \rightarrow G$ be an irreducible train track map that represents ϕ and let $\rho_G : R_r \rightarrow G$ be the marking on G . For all $i \geq 0$, let G_i be the marked graph obtained from $(g^i \rho_G, G)$ by multiplying the metric by $1/\lambda^i$, and so

$$\ell_{G_i}(c) = \ell_G(\phi^i(c))/\lambda^i \quad \text{for all } c \in \mathcal{C} \text{ and all } i \geq 0.$$

The map $g : G \rightarrow G$ induces a homotopy equivalence denoted $g_{i+1,i} : G_i \rightarrow G_{i+1}$ that preserves marking. By composition, for $0 \leq i < j$, the map $g^{j-i} : G \rightarrow G$ induces the marking preserving homotopy equivalence $g_{ji} = g_{j,j-1} \circ \cdots \circ g_{i+1,i} : G_i \rightarrow G_j$ (see Figure 1). We fix the base point of G_i so that the marking $g^i \rho_G : R_r \rightarrow G_i$ preserves base point. It follows that the maps $g_{ji} : G_i \rightarrow G_j$ preserve base point.

The sequence ℓ_{G_i} converges in $\widehat{\mathcal{X}}_r$ to an element $\ell_+ \in \widehat{\mathcal{X}}_r$ whose projective class is ϕ -invariant. To describe ℓ_+ explicitly, consider $c \in \mathcal{C}$ represented by a circuit $\sigma \subset G$. If $K(\sigma)$ is as in Lemma 2.7 and $k > K(\sigma)$ then $g_{\#}^k(\sigma)$ splits as a concatenation of legal subpaths $\mu_{i,k}$ and periodic indivisible Nielsen paths $\nu_{i,k}$. Excluding the $\nu_{i,k}$, define $\ell_+(c) = (\sum L_G(\mu_{i,k}))/\lambda^k$. Since g expands length of legal paths by the factor λ and there are only finitely many periodic indivisible Nielsen paths, $\ell_+(c)$ is independent of $k > K(\sigma)$ and is the limit of the sequence $\ell_{G_i}(c)$. It is immediate from the definition that $\ell_+(\phi(c)) = \lambda \ell_+(c)$ and that if c is represented by a legal circuit in G then $\ell_+(c) = \ell_G(c)$.

The F_r tree T_+ that realizes ℓ_+ is well defined up to isometric conjugacy.




Figure 1: Notation associated to a train track map $g: G \rightarrow G$.

Direct limits. It seems to be well known that T_+ is the F_r -equivariant Gromov limit of the trees \tilde{G}_i ; see e.g. [BF]. We need some slightly stronger facts about this convergence.

Direct limits exist in the category of metric spaces and distance nonincreasing maps. That is, consider a direct system consisting of a partially ordered set I , an indexed collection of metric spaces $(A_i)_{i \in I}$ with metrics d_i , and an indexed collection of distance nonincreasing maps $(f_{ji}: A_i \rightarrow A_j)_{i < j}$, such that I is a *directed set* meaning that for all $i, j \in I$ there exists $k \in I$ with $i \leq k, j \leq k$, and the maps satisfy the compatibility property $f_{kj} \circ f_{ji} = f_{ki}$ for $i < j < k$ in I . Then there exists a metric space A_∞ and distance nonincreasing maps $f_{\infty,i}: A_i \rightarrow A_\infty$ satisfying the following universality property: given a metric space B and a system of distance nonincreasing maps $h_i: A_i \rightarrow B$ compatible with the maps f_{ji} , meaning that $h_j \circ f_{ji} = h_i$, there exists a unique distance nonincreasing map $h_\infty: A_\infty \rightarrow B$ such that $h_\infty \circ f_{\infty,i} = h_i: A_i \rightarrow B$ for all $i \in I$. To construct A_∞ , first let \hat{A}_∞ be the direct limit of the system $(A_i), (f_{ji})$ in the category of sets [Spa81]: points of A_∞ are equivalence classes of the disjoint union of the A_i where $x \in A_i$ is equivalent to $y \in A_j$ if there exists k with $i, j \leq k$ such that $f_{ki}(x) = f_{kj}(y)$. Then define a pseudometric on \hat{A}_∞ where for any two equivalence classes $[x], [y]$ we have

$$d_\infty([x], [y]) = \inf_{k \geq i, j} d_k(f_{ki}(x), f_{kj}(y))$$

The direct limit A_∞ is the natural quotient metric space of \hat{A}_∞ , where $[x], [y] \in A_\infty$ are equivalent if and only if $d_\infty([x], [y]) = 0$. The map $f_{\infty,i}: A_i \rightarrow A_\infty$ is the composition of the quotient maps $A_i \rightarrow \hat{A}_\infty \rightarrow A_\infty$.

Consider the trees \tilde{G}_i . We lift the base point of G_i to \tilde{G}_i , and we choose lifts $\tilde{g}_{i+1,i}: \tilde{G}_i \rightarrow \tilde{G}_{i+1}$ of $g_{i+1,i}: G_i \rightarrow G_{i+1}$ for all i , so that the maps $\tilde{g}_{i+1,i}$ preserve base points, thereby making \tilde{G}_i into an F_r tree, and making each map $\tilde{g}_{i+1,i}$ into an F_r equivariant map which is an isometry when restricted to each edge. We obtain a direct system with $\tilde{g}_{ji}: \tilde{G}_i \rightarrow \tilde{G}_j$ defined inductively by $\tilde{g}_{ji} = \tilde{g}_{j,j-1} \circ \tilde{g}_{j-1,i}$. Whenever $\Phi \in \text{Aut}(F_r)$ representing ϕ is implicitly chosen, the associated lifts of $g_i: G_i \rightarrow G_i$ shall be denoted $\check{g}_i: \tilde{G}_i \rightarrow \tilde{G}_i$ for all $i \geq 0$ (see Figure 2). These maps satisfy the following properties:

- \check{g}_i restricted to each edge is a homothety with stretch equal to λ .
- $\tilde{g}_{ji}\check{g}_i = \check{g}_j\tilde{g}_{ji}$.
- (Twisted equivariance, see (2.1)) $\check{g}_i \circ t_\gamma = t_{\Phi(\gamma)} \circ \check{g}_i$ for each $\gamma \in F_r$.

While we recall that the notation \check{g} is meant to suggest twisted equivariance (2.1), on the other hand the notation \tilde{g} is meant to suggest F_r -equivariance.




Figure 2: Notation associated to a train track map $g: G \rightarrow G$ and automorphism Φ representing ϕ . Maps denoted \tilde{g} depend only on g and are F_r -equivariant. Maps denoted \check{g} depend on g and on Φ and satisfy the twisted equivariance equation (2.1). The maps f_{g_i} are defined after Theorem 2.13, and Φ_+ is defined in Theorem 2.15.

The following theorem and corollary seem to be folklore. See e.g. [HM06b] for a proof of Theorem 2.13, and see Section 7.3 below for a proof of a more general fact in the nongeometric case.

Theorem 2.13. *The direct limit of the direct system $\tilde{g}_{ij}: \tilde{G}_j \rightarrow \tilde{G}_i$, in the category of metric spaces with an F_r action whose morphisms are F_r equivariant, distance nonincreasing maps, is an F_r tree in the same homothety class as T_+ .* \diamond

In the context of this theorem, we shall normalize the metric on T_+ so that T_+ is identified with the direct limit. From Theorem 2.13 we obtain a direct limit map $f_g: \tilde{G} \rightarrow T_+$, and by construction f_g is an isometry when restricted to each edge of \tilde{G} . As the notation f_g suggests, this map is well-defined up to isometric conjugacy of the F_r action on T_+ (see Figure 2), depending only on $g: G \rightarrow G$, and *not* depending on the other choices made in the definition of the direct system $\tilde{g}_{ij}: \tilde{G}_j \rightarrow \tilde{G}_i$. This well-definedness is a consequence of the direct limit description of T_+ , from which it follows that T_+ is the metric space associated to a pseudometric on \tilde{G}_0 where the distance between $\tilde{x}, \tilde{y} \in \tilde{G}_0$ equals the limit of the nonincreasing sequence $\text{Len}\left(g_{\#}^k(\text{pr}[\tilde{x}, \tilde{y}])\right) / \lambda^k$. Clearly the pseudometric depends only on g and therefore so does T_+ .

Corollary 2.14. *Assume that $g: G \rightarrow G$ is rotationless and that \tilde{G} is a simplicial F_r -tree obtained from the universal cover of G by choosing a marking and a lifting of the base point. Then there is a surjective equivariant map $f_g: \tilde{G} \rightarrow T_+$ such that for all $\tilde{x}, \tilde{y} \in \tilde{G}$, letting $\sigma = \text{pr}[\tilde{x}, \tilde{y}]$, the following are equivalent.*

- (1) $f_g(\tilde{x}) = f_g(\tilde{y})$.
- (2) For some $k \geq 0$, the path $g_{\#}^k(\sigma)$ is either trivial or Nielsen.

In particular, f_g restricts to an isometry on all legal paths.

Proof. Assume the notation of Theorem 2.13 with $\tilde{G} = \tilde{G}_0$ and let $f_g: \tilde{G} \rightarrow T_+$ be the direct limit map.

We first show that (2) implies (1). By construction, $g_{\#}^k(\sigma)$ is trivial if and only if $\tilde{g}_{k,0}(\tilde{x}) = \tilde{g}_{k,0}(\tilde{y})$ and the latter implies that $f_g(\tilde{x}) = f_g(\tilde{y})$. To prove (1) assuming that σ is a Nielsen or pre-Nielsen path, it suffices to assume that σ is indivisible. By Lemma 2.7-(2), there exist $\tilde{x}_i \rightarrow \tilde{x}$ and $\tilde{y}_i \rightarrow \tilde{y}$ such that $\tilde{g}_{k,0}(\tilde{x}_i) = \tilde{g}_{k,0}(\tilde{y}_i)$ for all i and all sufficiently large k . It follows that $f_g(\tilde{x}_i) = f_g(\tilde{y}_i)$ for all i and hence that $f_g(\tilde{x}) = f_g(\tilde{y})$.

Suppose now that $f_g(\tilde{x}) = f_g(\tilde{y})$. If each $g_{\#}^k(\sigma)$ is non-trivial, then Lemma 2.7 implies that $g_{\#}^k(\sigma)$ splits as a concatenation of legal paths and indivisible Nielsen paths for all sufficiently large k . If μ is any legal path in this splitting, then the distance between $\tilde{g}_{k,0}(\tilde{x})$ and $\tilde{g}_{k,0}(\tilde{y})$ is uniformly bounded below by the length of μ in contradiction to the definition of f_g and the assumption that $f_g(\tilde{x}) = f_g(\tilde{y})$. Thus no such μ exists and $g_{\#}^k(\sigma)$ is Nielsen.

We have now proven the equivalence of (1) and (2). Since a legal path does not contain subpaths σ satisfying (2), the restriction of \tilde{g}_{ij} to any legal path is an embedding, and hence an isometry, and it follows that f_g restricts to an isometry on all legal paths. \diamond

Next we review the well known fact that there is a kind of action on the \mathbf{R} -tree T_+^ϕ by the set of automorphisms of F_r representing a given fully irreducible outer automorphism ϕ . This set is not a group, instead it is a coset of the normal subgroup $\text{Inn}(F_r)$ in the group $\text{Aut}(F_r)$ — formally this coset is identified with the outer automorphism $\phi \in \text{Out}(F_r) = \text{Aut}(F_r)/\text{Inn}(F_r)$. So while it does not make sense to ask for a “group action” of the coset, it does make sense to ask for an “ $\text{Inn}(F_r)$ -affine action”, which is what item (2) gives us:

Theorem 2.15. *Let $\phi \in \text{Out}(F_r)$ be fully irreducible, and denote $T_+ = T_+^\phi$. To each $\Phi \in \text{Aut}(F_r)$ representing ϕ there is associated a homothety $\Phi_+ : T_+ \rightarrow T_+$ with stretch factor λ satisfying the following properties:*

- (1) *For each irreducible train track representative $g : G \rightarrow G$ of ϕ , letting $\check{g} : \tilde{G} \rightarrow \tilde{G}$ be the lift corresponding to Φ , we have*

$$\Phi_+ \circ f_g = f_g \circ \check{g}$$

- (2) *For each Φ, Φ' representing ϕ , if $\gamma \in F_r$ is the unique element such that $\Phi' = i_\gamma \circ \Phi$ then $\Phi'_+ = t_\gamma \circ \Phi_+$.*

Proof. Choose any train track representative $g : G \rightarrow G$ of ϕ , and refer to Figure 1 for notations associated to g . Let $\check{g}, \check{g}' : \tilde{G} \rightarrow \tilde{G}$ be the lifts associated to Φ, Φ' , respectively. Refer to Figure 2 for notations associated to \check{g} . Refer again to Figure 2 for notations associated to \check{g}' , but with the prime symbol $'$ added to these notations. In particular we have identifications $\tilde{G} = \tilde{G}_0$, $\check{g} = \check{g}_0$, $\check{g}' = \check{g}'_0$.

Let T_+^* be T_+ with the metric scaled down by the factor λ and let $h : T_+ \rightarrow T_+^*$ be the ‘identity’ map; thus h contracts length uniformly by λ . By the universal property of direct limits applied to the maps $h \circ f_{g_i} \circ \check{g}_i : \tilde{G}_i \rightarrow T_+^*$ there is a map $\Phi_+^* : T_+ \rightarrow T_+^*$ such that $h \circ f_{g_i} \circ \check{g}_i = \Phi_+^* \circ f_{g_i}$. Let $\Phi_+ : T_+ \rightarrow T_+$ be $h^{-1} \circ \Phi_+^*$. It follows that $f_{g_i} \circ \check{g}_i = \Phi_+ \circ f_{g_i}$, which implies (1) by taking $i = 0$. Similarly Φ'_+ is defined satisfying $f_{g_i} \circ \check{g}'_i = \Phi'_+ \circ f_{g_i}$, which implies (2) together with the fact that $\check{g}'_i = t_\gamma \circ \check{g}_i$ and with the universal property of direct limits.

Given z_1 and z_2 in T_+ , choose $\tilde{x}_j \in \tilde{G}_0$ such that $f_0(\tilde{x}_j) = z_j$. By the third item of Lemma 2.7, the projection into G_i of the path in \tilde{G}_i connecting $\tilde{g}_{i0}(x_1)$ to $\tilde{g}_{i0}(x_2)$ splits

into subpaths that are either legal or indivisible Nielsen paths for all sufficiently large i . Corollary 2.14 implies that the distance between z_1 and z_2 is the sum of the lengths of the legal subpaths of G_i in this splitting. Since g expands length by a factor of λ on all legal subpaths, (1) implies that Φ_+ must expand the distance between z_1 and z_2 by this same factor and since z_1 and z_2 were arbitrary, it follows that Φ_+ expands length uniformly by λ . This proves that Φ_+ is a homothety onto its image with stretch factor λ , but Φ_+ is surjective by minimality of the F_r -tree T_+ . \diamond

Next we describe the connection between principal automorphisms representing a fully irreducible $\phi \in \text{Out}(F_r)$ and the set of branch points of T_+^ϕ .

Lemma 2.16. *Suppose that $\phi \in \text{Out}(F_r)$ is fully irreducible and rotationless, $g: G \rightarrow G$ is a train track representative, Φ is a principal automorphism representing ϕ and $\check{g}: \tilde{G} \rightarrow \tilde{G}$ is the principle lift corresponding to Φ , and $\Phi_+: T_+ \rightarrow T_+$ is as in Theorem 2.15. Then*

- (1) $f_g(\text{Fix}(\check{g}))$ is a branch point b of T_+ , and $\text{Fix}(\Phi_+) = \{b\}$.
- (2) Every direction based at b is fixed by $D\Phi_+$ and has the form $Df_g(d)$ where d is a fixed direction based at some $\tilde{v} \in \text{Fix}(\check{g})$.
- (3) If $f_g(\tilde{x}) = b$ and d is a direction based at \tilde{x} then
 - For some $m > 0$, $\check{g}^m(\tilde{x}) \in \text{Fix}(\check{g})$.
 - $f_g(\check{g}^k(\tilde{x}))$ and $Df_g(D\check{g}^k(d))$ are independent of $k \geq 0$.
- (4) The assignment $\Phi \mapsto b = \text{Fix}(\Phi_+)$ of (1) defines a bijection between $\text{PA}(\phi)$ and the set of branch points of T_+ .

Proof. First we prove (1). Lemma 2.11 implies that $\text{Fix}(\check{g}) \neq \emptyset$ and that the elements of $\text{Fix}(\check{g})$ are principal vertices. By Corollary 2.14 and Theorem 2.15, the point $b = f_g(\tilde{v})$ is independent of $\tilde{v} \in \text{Fix}(\check{g})$ and $\{b\} = \text{Fix}(\Phi_+)$, and there is an induced map Df_g from directions at \tilde{v} to directions at b that is injective on the set of fixed directions. If there are at least three fixed directions at \tilde{v} then b is a branch point. Otherwise, there is an indivisible Nielsen path σ and a lift $\tilde{\sigma}$ with one endpoint at \tilde{v} and the other at some principal $\tilde{w} \in \text{Fix}(\check{g})$. Let d_1 and d_2 be distinct fixed directions at \tilde{v} with d_2 being the initial direction of $\tilde{\sigma}$ and let d_3 and d_4 be distinct fixed directions at \tilde{w} with d_3 being the initial direction of the inverse of $\tilde{\sigma}$. Then $Df_g(d_1) \neq Df_g(d_2) = Df_g(d_3) \neq Df_g(d_4)$. To prove that b is a branch point, it suffices to prove that $Df_g(d_1) \neq Df_g(d_4)$. If this fails then there exist \tilde{x} close to \tilde{v} in the direction d_1 and \tilde{y} close to \tilde{w} in the direction d_4 such that $f_g(\tilde{x}) = f_g(\tilde{y})$. The second item of Lemma 2.7 implies that the projection into G of the path connecting \tilde{x} to \tilde{y} splits as a non-trivial legal path followed by the indivisible Nielsen path σ followed by a non-trivial legal path. Corollary 2.14 then implies that \tilde{x} and \tilde{y} have distinct \tilde{f}_g -images and this contradiction completes the proof of item (1).

We now turn to (3). Choose $\tilde{v} \in \text{Fix}(\check{g})$. Corollary 2.14 implies that if $f_g(\tilde{x}) = b$ then there exists $m \geq 0$ such that either $\check{g}^m(\tilde{x}) = \tilde{v}$ or the path connecting $\check{g}^m(\tilde{x})$ to \tilde{v} is the lift of a Nielsen path. In either case, $\check{g}^m(\tilde{x}) \in \text{Fix}(\check{g})$, which implies that $f_g(\check{g}^k(\tilde{x})) = b$ for all $k \geq 0$. After increasing m , we may assume that $D\check{g}^m(d) \in \text{Fix}(D\check{g})$ for all directions d based at \tilde{x} . From the fact that $f_g \circ \check{g} = \Phi_+ \circ f_g$ we conclude that $D\Phi_+^k \circ Df_g(d) = Df_g(D\check{g}^k(d))$ is

independent of $k \geq m$. Since Φ_+ is a homeomorphism, they are in fact independent of $k \geq 0$. This completes the proof of item (3) and also implies that $Df_g(d) \in \text{Fix}(D\Phi_+)$. Since every direction at b occurs as $Df_g(d)$ for some \tilde{x} and d , we have also proved item (2).

Next we prove injectivity of the map $\Phi \mapsto b$. Suppose that \tilde{v}_1 and \tilde{v}_2 are distinct principal vertices and that $f_g(\tilde{v}_1) = f_g(\tilde{v}_2)$. Let \check{g}_1 be the lift that fixes \tilde{v}_1 . Since principal vertices of \tilde{G} project into $\text{Fix}(g)$, the restriction of \check{g}_1 to the set of principal vertices of \tilde{G} is a bijection. It follows that $\check{g}^k(\tilde{v}_1) \neq \check{g}^k(\tilde{v}_2)$ for all k . The path $\tilde{\sigma}$ connecting \tilde{v}_1 to \tilde{v}_2 cannot project to a pre-Nielsen path σ which is not Nielsen. Indeed if it did, then for some $k > 0$, the Nielsen path $\sigma' = f_{\#}^k(\sigma)$ has the same endpoints as σ and $f_{\#}^k(\sigma') = f_{\#}^k(\sigma)$ which contradicts that $\sigma' \neq \sigma$ and the fact that f is a homotopy equivalence. Corollary 2.14 therefore implies that $\tilde{\sigma}$ projects to a Nielsen path and hence that $\tilde{v}_2 \in \text{Fix}(\check{g}_1)$. This proves that $\Phi \mapsto b$ is injective.

To see that $\Phi \mapsto b$ is onto it suffices to show that if b is a branch point then $f_g(\tilde{c}) = b$ for some principal vertex \tilde{c} ; one then chooses Φ to be the automorphism corresponding to the lift of g that fixes \tilde{c} . For $1 \leq l \leq 3$ there are points $\tilde{a}_l \in \tilde{G}$ and directions d_l at \tilde{a}_l such that $f_g(\tilde{a}_l) = b$, and such that $Df_g(d_l)$ are three distinct directions at b . Let $a_j \in G$ be the projected image of \tilde{a}_j . On the one hand, if $g^m(a_1) = g^m(a_2) = g^m(a_3)$ for some $m \geq 0$ then there are three directions at this point that have distinct Df_g -images and so must have distinct Dg^k -images for all $k \geq 0$. Otherwise, Corollary 2.14 implies that $g^m(a_j)$ is the endpoint of a Nielsen path for some j and some $m \geq 0$. In either case $c := g^m(a_j)$ is a principal, and hence fixed, vertex. There is a path σ connecting a_j to c so that $g_{\#}^m(\sigma)$ is trivial. The lift $\tilde{\sigma}$ with endpoint \tilde{a}_j terminates at a principal vertex \tilde{c} and $\check{g}(\tilde{c}) = \check{g}(\tilde{a}_j) = b$. This completes the proof of (4) and so the proof of the lemma. \diamond

2.6 Geodesic laminations in trees and marked graphs

Given a topological space S let D^2S denote the double space of S , the space of distinct ordered pairs in S :

$$D^2S = \{(\xi, \eta) \in S \times S \mid \xi \neq \eta\}$$

The group $\mathbf{Z}/2$ acts freely on D^2S by permuting coordinates.

Define the *geodesic leaf space* of F_r , denoted $\mathcal{G}F_r$, to be the set of distinct *unordered* pairs of points in ∂F_r , that is, the quotient space of the action of $\mathbf{Z}/2$ on $D^2\partial F_r$:

$$\mathcal{G}F_r = (D^2\partial F_r) / \mathbf{Z}/2$$

The space $\mathcal{G}F_r$ is locally a Cantor set. We will usually abuse ordered pair notation and write (ξ, η) for a point in $\mathcal{G}F_r$. The action of F_r on ∂F_r induces an action on $\mathcal{G}F_r$ which has a dense orbit. The quotient space $\mathcal{G}F_r/F_r$ is therefore a non-Hausdorff topological space; this space is denoted \mathcal{B} in [BFH00].

Given a free, simplicial F_r -tree T , we also define the geodesic leaf space

$$\mathcal{G}T = (D^2\partial T) / (\mathbf{Z}/2)$$

The natural equivariant homeomorphism $\partial F_r \approx \partial T$ induces a natural equivariant homeomorphism $\mathcal{G}F_r \approx \mathcal{G}T$, and we will henceforth identify $\mathcal{G}T$ with $\mathcal{G}F_r$.

It will be convenient to think of \mathcal{GT} not just in the abstract as a set of unordered distinct pairs of boundary points, but more concretely as a set whose elements are bi-infinite geodesics, and even more concretely as a lamination in the sense of topology or dynamics, which is a topological space decomposed into 1-manifolds fitting together locally as an interval crossed with a transversal. This approach is used, for example, in [CP94]. One advantage of this approach is that it avoids conceptual difficulties of non-Hausdorff spaces. We briefly review this point of view.

Define a *parameterized geodesic* in T to be an isometric embedding $\gamma: \mathbf{R} \rightarrow T$, and define the *geodesic flow* of T , denoted \mathcal{GFT} , to be the space of parameterized geodesics with the compact open topology. The group \mathbf{R} acts homeomorphically on \mathcal{GFT} by translating the domain: $\gamma \cdot r(t) = \gamma(t + r)$. An orbit of the \mathbf{R} action on \mathcal{GFT} is called a *geodesic flow line*, and there are natural bijections between the set of geodesic flow lines, the set of oriented bi-infinite geodesics, and the set $D^2\partial T$. The geodesic flow lines decompose \mathcal{GFT} , and they fit together locally like the product of an interval crossed with a Cantor set, that is, a 1-dimensional lamination with Cantor cross sections. The group $\mathbf{Z}/2$ acts on \mathcal{GFT} by reflecting the domain \mathbf{R} through the origin. This action is free on the set of geodesic flow lines, interchanging a flow line with its orientation reversal. The quotient $\mathcal{GLT} = \mathcal{GFT}/(\mathbf{Z}/2)$ is called the *geodesic lamination* of T . An element of \mathcal{GLT} can be naturally identified with a “pointed geodesic” in T , a bi-infinite geodesic equipped with a base point. The decomposition of \mathcal{GFT} into flow lines descends to a decomposition of \mathcal{GLT} into sets called *geodesic leaves*, where each leaf corresponds to different choices of base points along some bi-infinite geodesic in T ; the geodesic leaves form a 1-dimensional lamination with Cantor cross sections. We obtain natural bijections between the set of geodesic leaves, the set of unoriented bi-infinite geodesics, and the set \mathcal{GT} .

The action of F_r on \mathcal{GT} is free, properly discontinuous, and cocompact, and the quotient space \mathcal{GT}/F_r inherits the structure of a 1-dimensional lamination with Cantor cross sections. We can view this more intrinsically as the *geodesic lamination* of the quotient marked graph $G = T/F_r$, as follows. We can define the geodesic flow \mathcal{GFG} as above, as parameterized local geodesics $\gamma: \mathbf{R} \rightarrow G$ with the compact open topology, equipped with the \mathbf{R} action by translation of the domain. The group $\mathbf{Z}/2$ acts by reflecting the domain through the origin, and we obtain the quotient $\mathcal{GLG} = \mathcal{GFG}/(\mathbf{Z}/2)$. Both \mathcal{GFG} and \mathcal{GLG} are 1-dimensional laminations with Cantor cross sections, and there are natural identifications $\mathcal{GFG} \approx \mathcal{GFT}/F_r$, $\mathcal{GLG} = \mathcal{GLT}/F_r$.

Given $\lambda = (\xi, \eta) \in \mathcal{GF}_r$ let λ_T denote the leaf of \mathcal{GLT} with ideal endpoints ξ, η ; we say that λ_T is the *realization* of $\lambda \in \mathcal{GF}_r$. This relation is an equivariant bijection between points of \mathcal{GF}_r and leaves of \mathcal{GLT} . Moreover this relation respects the topologies in the following sense: a sequence $\lambda_i \in \mathcal{GF}_r$ converges to $\lambda \in \mathcal{GF}_r$ if and only if the sequence $(\lambda_i)_T$ converges to λ_T in the Gromov-Hausdorff topology on closed subsets of T , meaning that the intersections with any closed ball in T converge in the Hausdorff topology on closed subsets of that ball.

A *sublamination* of \mathcal{GF}_r is a closed, F_r -equivariant subset of \mathcal{GF}_r . Given a free, simplicial F_r -tree T , a *sublamination* of \mathcal{GLT} is a closed, F_r -equivariant union of leaves of \mathcal{GLT} . And, given a marked graph G , a *sublamination* of \mathcal{GLG} is a closed union of leaves of \mathcal{GLG} . For any T and for $G = T/F_r$, and for any sublamination $\Lambda \subset \mathcal{GF}_r$, the *realization* of Λ in T is

the sublamination of $\mathcal{GL}T$ defined by

$$\Lambda_T = \{\lambda_T \mid \lambda \in \Lambda\}$$

Also, the *realization* of Λ in G is the quotient $\Lambda_G = \Lambda_T/F_r$, a sublamination of $\mathcal{GL}G$. This gives bijections between sublaminations of \mathcal{GF}_r , sublaminations of $\mathcal{GL}T$, and sublaminations of $\mathcal{GL}G$. We make free use of these bijections without comment, and we freely use the symbol Λ for sublaminations in any of these contexts. When Λ stands for a sublamination of \mathcal{GF}_r then we sometimes use $\Lambda(T)$ or $\Lambda(G)$ instead of Λ_T or Λ_G . When Λ stands for a sublamination of $\mathcal{GL}G$, the corresponding sublamination of $\mathcal{GL}T$ is usually denoted $\tilde{\Lambda}$.

A sublamination $\Lambda \subset \mathcal{GF}_r$ is *minimal* if and only if every F_r -orbit in Λ is dense in Λ ; equivalently, for each free simplicial F_r -tree T , each orbit of leaves of Λ_T is dense in Λ_T ; equivalently, for each marked graph G , each leaf of Λ_G is dense in Λ_G .

The group $\text{Aut}(F_r)$ acts on ∂F_r and therefore also on the geodesic leaf space \mathcal{GF}_r . Note that this action agrees, under restriction, with the given action of $F_r \approx \text{Inn}(F_r)$ on \mathcal{GF}_r . Choosing a marked graph G with universal cover T , \mathcal{GF}_r is identified with the set of leaves of $\mathcal{GL}T$ and so $\text{Aut}(F_r)$ acts on the latter, in agreement under restriction with the action of $F_r \approx \text{Inn}(F_r)$. Modding out by $F_r \approx \text{Inn}(F_r)$ we obtain an action of $\text{Out}(F_r)$ on the set of leaves of $\mathcal{GL}G$. We will use these actions without reference in what follows.

2.7 The expanding lamination Λ_-

The chief example of a sublamination of \mathcal{GF}_r is the expanding lamination Λ_- of a fully irreducible outer automorphism $\phi \in \text{Out}(F_r)$, as constructed in [BFH97]. We review here the construction, well-definedness, and properties of Λ_- . In particular we present key features of Λ_- in Lemmas 2.17 and 2.18, the second of which motivates our choice of notation. The construction of Λ_- will be revisited in Section 7.4 when we extend the construction to a more general context. Complete details on Λ_- can be found in either section 1 of [BFH97] or section 3 of [BFH00].

Definition and basic properties of Λ_- . Consider any train track representative $g: G \rightarrow G$ of ϕ . Choose a sequence of pairs (E_i, x_i) so that each E_i is an edge of G , $x_i \in \text{int}(E_i)$, and we assume that $\lambda^i d(x_i, \partial E_i) \rightarrow +\infty$; this is true, for example, if (x_i) is a periodic orbit. Then $(g^i(E_i), g^i(x_i))$ is a sequence of pointed geodesic arcs whose lengths go to $+\infty$ so that the distance from the base point to the endpoints also goes to $+\infty$. We may therefore choose a subsequence of (E_i, x_i) so that $(g^i(E_i), g^i(x_i))$ converges to a bi-infinite, pointed geodesic in G . The union of such geodesics, over all the choices made, is clearly a sublamination of $\mathcal{GL}G$ that we denote $\Lambda_-(g)$. In [BFH97] it is proved that the expanding lamination $\Lambda_- \subset \mathcal{GF}_r$ of ϕ is well-defined and that its realization in G , for each train track representative $g: G \rightarrow G$ of ϕ , is $\Lambda_-(g)$. If the train track map is clear from the context then we will simply denote $\Lambda_-(g)$ as $\Lambda_-(\phi)$ or just as Λ_- .

The expanding lamination for ϕ^p , $p > 0$, is evidently the same as the expanding lamination for ϕ so we may replace ϕ by an iterate when it is convenient.

By irreducibility of $g: G \rightarrow G$, every leaf or half-leaf of $\Lambda_-(g)$ covers every edge of G . Also, Λ_- is carried by no proper free factor of F_r (see Example 2.5 (1) in [BFH97]), so the realization of Λ_- in any marked graph has the property that every leaf covers every edge.

Action of ϕ on Λ_- . Under the action of $\text{Out}(F_r)$ on the set of leaves of \mathcal{GLG} described at the end of Section 2.6, the outer automorphism ϕ leaves invariant the set of leaves of Λ_- . We denote this restricted action by $\ell \mapsto \phi_{\#}(\ell)$, for each leaf ℓ of Λ_- . This action is characterized by saying that for each train track representative $g: G \rightarrow G$ of ϕ and each leaf ℓ of Λ_- , g takes the realization of ℓ in G to the realization of $\phi_{\#}(\ell)$ in G . For any $\Phi \in \text{Aut}(F_r)$ that represents ϕ , there is an action of Φ on the set of leaves of $\tilde{\Lambda}_-$, denoted $\tilde{\ell} \mapsto \Phi_{\#}(\tilde{\ell})$, which is well-defined by saying that for each train track representative $g: G \rightarrow G$, letting $\tilde{g}: \tilde{G} \rightarrow \tilde{G}$ be the lift of g corresponding to Φ , the map \tilde{g} takes the realization of $\tilde{\ell}$ in \tilde{G} to the realization of $\Phi_{\#}(\tilde{\ell})$ in \tilde{G} . To justify the existence of these actions see the final comments of Section 2.6.

A leaf ℓ of Λ_- is *periodic* if $\phi_{\#}^i(\ell) = \ell$ for some $i \geq 1$, and a leaf $\tilde{\ell}$ of $\tilde{\Lambda}_-$ is periodic if there exists $\Phi \in \text{Aut}(F_r)$ representing an iterate of ϕ such that $\Phi_{\#}(\tilde{\ell}) = \tilde{\ell}$. Periodicity of leaves is preserved under projection from $\tilde{\Lambda}_-$ to Λ_- .

Λ_- and local Whitehead graphs. From the definition of $\Lambda_-(g)$, for any $x \in G$ the local Whitehead graph $W(x; G)$ has a description in terms of $\Lambda_-(g)$, as follows. A nondegenerate turn (d_1, d_2) at x is taken by some leaf of $\Lambda_-(g)$ if and only if this turn is taken by $g^k(E)$ for some edge E of G and some $k \geq 1$, which holds if and only if (d_1, d_2) are the endpoints of an edge of the local Whitehead graph $W(x; G)$.

Singular leaves of Λ_- and stable Whitehead graphs. Next we introduce notation and terminology to refer to certain leaves of Λ_- , called “singular leaves” in analogy to singular leaves of the unstable foliation of a pseudo-Anosov surface homeomorphism. These singular leaves can be used to relate Λ_- to stable Whitehead graphs.

Let $g: G \rightarrow G$ be an irreducible train track map representing an iterate of ϕ such each principal or nonprincipal periodic vertex v is fixed and each periodic direction at v is fixed. Choose a lift $\tilde{v} \in \tilde{G}$ and a lift $\tilde{g}: \tilde{G} \rightarrow \tilde{G}$ that fixes \tilde{v} . If d is a direction at \tilde{v} fixed by $D\tilde{g}$ and \tilde{E} is the edge at \tilde{v} corresponding to d , then $\tilde{g}^j(\tilde{E})$ is a proper initial subpath of $\tilde{g}^{j+1}(\tilde{E})$ for each $j \geq 0$ and so $\tilde{R} = \bigcup_{j=0}^{\infty} \tilde{g}^j(\tilde{E})$ is a ray in \tilde{G} converging to some $P \in \text{Fix}_N(\hat{g})$. We say that \tilde{R} is the *ray determined by d* (or *by \tilde{E}*), and if v is principal then \tilde{R} is a *singular ray*.

Let \tilde{R}_1, \tilde{R}_2 be rays at \tilde{v} determined by distinct fixed directions d_1, d_2 , corresponding to edges \tilde{E}_1, \tilde{E}_2 at \tilde{v} . We claim that if d_1 and d_2 are endpoints of an edge in $SW(\tilde{v}; \tilde{G})$ (which is always true if \tilde{v} is not principal), then $\tilde{\ell} = \tilde{R}_1 \cup \tilde{R}_2$ is a leaf of $\tilde{\Lambda}_-$ realized in \tilde{G} . This follows because there exists an edge \tilde{E} of \tilde{G} and $l > 0$ such that $\tilde{g}^l(\tilde{E})$ contains the turn (d_1, d_2) , and so as $i \rightarrow +\infty$ the paths $\tilde{g}^{l+i}(\tilde{E})$ converge (in the Gromov–Hausdorff topology on closed subsets) to $\tilde{\ell} = \tilde{R}_1 \cup \tilde{R}_2$. The leaf $\tilde{\ell}$ is called *the leaf at \tilde{v} determined by d_1 and d_2* (or by \tilde{E}_1 and \tilde{E}_2). If \tilde{v} is principal then we say that $\tilde{\ell}$ is a *singular leaf*.

In the previous paragraph, the converse to the claim holds as well: if $\tilde{\ell} = \tilde{R}_1 \cup \tilde{R}_2$ is a leaf of $\tilde{\Lambda}_-$ realized in \tilde{G} then d_1, d_2 bound an edge of $SW(\tilde{v}; \tilde{G})$. This holds because, by definition of $\tilde{\Lambda}_-$, each segment of $\tilde{\ell}$ must be contained in $\tilde{g}^l(\tilde{E})$ for some edge \tilde{E} and some $l > 0$. It follows that $SW(\tilde{v}; \tilde{G})$ can be identified with the graph having one vertex for each singular ray \tilde{R} based at \tilde{v} and having an edge connecting the vertices corresponding to a pair of singular rays \tilde{R}_1, \tilde{R}_2 at \tilde{v} if and only if $\tilde{R}_1 \cup \tilde{R}_2$ is a singular leaf at \tilde{v} .

Characterizing Λ_- . The next lemma gives a useful characterization of Λ_- , and its corollary characterizes the nonsingular ϕ -periodic leaves of Λ_- .

Lemma 2.17. *Λ_- is the unique minimal ϕ -invariant lamination whose leaves are legal with respect to every train track representative $g: G \rightarrow G$ of ϕ . Every leaf $\tilde{\ell}$ constructed as above, and in particular every singular leaf, is contained in Λ_- .*

Proof. With notation as above, consider the set \mathcal{L} of minimal ϕ -invariant laminations whose leaves are legal with respect to every train track representative $g: G \rightarrow G$ of ϕ . That $\Lambda_- \in \mathcal{L}$ follows from the construction of Λ_- in section 1 of [BFH97], as reviewed above. We will show that if $\Lambda' \in \mathcal{L}$ and $\tilde{\ell}$ is any leaf constructed as above, then $\tilde{\ell}$ is a leaf of Λ' . Since there can only be one minimal lamination that contains $\tilde{\ell}$ it follows Λ_- is the only element of \mathcal{L} .

Pass to a power and choose an appropriate lift $\check{g}: \check{G} \rightarrow \check{G}$ as above, and choose $\tilde{v} \in \check{G}$ and distinct edges \tilde{E}_1 and \tilde{E}_2 at \tilde{v} all of which are fixed by \check{g} , so that \tilde{E}_1, \tilde{E}_2 determine an edge in $SW(\tilde{v}; \check{G})$. There is an edge \tilde{E} and $l > 0$ such that $\check{g}^l(\tilde{E})$ contains the turn $(\tilde{E}_1, \tilde{E}_2)$. Since g is irreducible and leaves of Λ' are legal in G , there is a leaf $\tilde{\ell}'$ of $\tilde{\Lambda}'$ that contains \tilde{E} . Each path in the nested increasing sequence $\tilde{E}_1^{-1}\tilde{E}_2 \subset \check{g}(\tilde{E}_1^{-1}\tilde{E}_2) \subset \check{g}^2(\tilde{E}_1^{-1}\tilde{E}_2) \subset \dots$ occurs as a subpath of a leaf of $\tilde{\Lambda}'$. It follows that their union $\tilde{\ell}$ is a leaf of $\tilde{\Lambda}'$. \diamond

Corollary 2.18. *If $\Phi \in \text{Aut}(F_n)$ represents an iterate of ϕ and $\text{Per}_N(\hat{\Phi}) = \text{Fix}_N(\hat{\Phi}) = \{P_1, P_2\}$ then the line $\tilde{\ell}$ bounded by P_1 and P_2 is a periodic nonsingular leaf of $\tilde{\Lambda}$. Conversely, if $\tilde{\ell}$ is a nonsingular periodic leaf of $\tilde{\Lambda}$ then there exists $\Phi \in \text{Aut}(F_n)$ representing an iterate of ϕ such that $\text{Per}_N(\hat{\Phi}) = \text{Fix}_N(\hat{\Phi}) = \partial\tilde{\ell}$.*

Proof. Choose an irreducible train track map $g: G \rightarrow G$ representing ϕ and let \check{g}_k be the lift of an iterate of g that corresponds to Φ . Since P_1 and P_2 are attractors, $\text{Fix}(\check{g}_k) \neq \emptyset$; see for example Lemma 3.9 and Lemma 3.13 of [FH04]. After replacing \check{g}_k by an iterate we may choose $\tilde{x} \in \text{Fix}(\check{g}_k)$ and a pair of fixed directions d_1 and d_2 at \tilde{x} . After subdividing G , we may assume that \tilde{x} is a vertex. If there were any other periodic direction at \tilde{x} then the corresponding singular ray would give another point in $\text{Per}_N(\hat{\Phi})$ in addition to P_1, P_2 , a contradiction. It follows that the stable Whitehead graph $SW(\tilde{x}; \check{G})$ has just the two vertices d_1, d_2 , and moreover the rays R_1, R_2 determined by d_1, d_2 must end at points of $\text{Per}_N(\hat{\Phi})$, and up to re-indexing these points must be P_1, P_2 . Since $SW(\tilde{x}; \check{G})$ is connected (see Section 2.4), there must be an edge of $SW(\tilde{x}; \check{G})$ connecting d_1 to d_2 . An application of Lemma 2.17 shows that $\tilde{\ell} = R_1 \cup R_2$ is in $\tilde{\Lambda}$.

For the converse, let $\tilde{\ell}$ be a nonsingular periodic leaf of $\tilde{\Lambda}$. Choose $\Phi \in \text{Aut}(F_r)$ representing an iterate of ϕ so that $\Phi\#(\tilde{\ell}) = \tilde{\ell}$. Choose an irreducible train track map $g: G \rightarrow G$ representing ϕ and let $\check{g}_k: \check{G} \rightarrow \check{G}$ be the lift of an iterate of g that corresponds to Φ . Using $\tilde{\ell}$ to denote its own realization in \check{G} , we have $\check{g}_k(\tilde{\ell}) = \tilde{\ell}$. Since \check{g}_k expands $\tilde{\ell}$ by the uniform factor λ , there is a point $x \in \tilde{\ell}$ fixed by \check{g}_k . Subdivide G so that x is a vertex. Up to passing to the square, $D\check{g}_k$ fixes the two directions d_1, d_2 of $\tilde{\ell}$ at x , and the rays determined by these two directions are the two rays of $\tilde{\ell}$ based at x . It follows that $\partial\tilde{\ell} \subset \text{Fix}_N(\hat{\Phi}) \subset \text{Per}_N(\hat{\Phi})$. If there were another point in $\text{Per}_N(\hat{\Phi})$ then, passing to a power so that this point is in $\text{Fix}_N(\hat{\Phi})$, it would follow that $\tilde{\ell}$ is singular, contrary to hypothesis. \diamond

2.8 Relating Λ_- to T_- and to T_+

Consider an \mathbf{R} -tree T representing a point in $\partial\mathcal{X}_r$, and a nonempty minimal sublamination Λ of \mathcal{GF}_r . Following [BFH97] section 3 we say that Λ has length zero in T if for every $G \in \mathcal{X}_r$ and every F_r -equivariant morphism $h: \tilde{G} \rightarrow T$ there exists $C \geq 0$ such that for every leaf $\ell \subset \tilde{G}$ of $\tilde{\Lambda}_G$ we have $\text{diam}_{T_-}(h(\ell)) \leq C$.

Lemma 2.19. *For each fully irreducible $\phi \in \text{Out}(F_r)$, $\Lambda_- = \Lambda_-(\phi)$ is the unique minimal lamination that has length zero in $T_- = T_-(\phi)$.*

Proof. This is proved in Lemma 3.5 (4) of [BFH97], but restricted to the so-called “irreducible laminations”, meaning the set of expanding laminations of all fully irreducible outer automorphisms in F_r . However, the proof applies to arbitrary minimal sublaminations, for as is noted several times in [BFH97], the only property of irreducibility of Λ that is really used is quasiperiodicity of leaves of Λ , which is equivalent to minimality of Λ . See for example the proof of [BFH97] Proposition 2.12 which is used in the proof of Lemma 3.5 (4). \diamond

Remark Lemma 2.19 shows that the expanding lamination of a fully irreducible outer automorphism ϕ is naturally related to the repelling tree T_- rather than the attracting tree T_+ . This is why we have chosen to denote the expanding lamination with a subscript “-”, altering the notation from the subscript “+” used in [BFH97].

We next relate Λ_- to T_+ .

Lemma 2.20. *For any train track representative $g: G \rightarrow G$ of ϕ and any leaf $\tilde{\ell}$ of $\tilde{\Lambda}_-$, the restriction of the map $f_g: \tilde{G} \rightarrow T_+$ to the realization of $\tilde{\ell}$ in \tilde{G} is an isometry. Moreover, the bi-infinite line in T_+ which is the image of this isometry is independent of the choice of g , depending only on $\tilde{\ell}$; this image is called the realization of $\tilde{\ell}$ in T_+ .* \diamond

Motivated by the first part of this lemma, given a free, simplicial F_r -tree T , an F_r -equivariant map $f: T \rightarrow T_+$ is said to be a Λ_- isometry, or to be Λ_- isometric, if the restriction of f to any leaf of $\tilde{\Lambda}_-$ realized in T is an isometry. The first part of the lemma therefore says that the map f_g is a Λ_- isometry.

Remark. The property that T has a Λ_- isometry $f: T \rightarrow T_+$ does not, in and of itself, determine the map f . In Section 5.4 we give an example of a marked graph G such that there exist at least two distinct Λ_- isometries $\tilde{G} \rightarrow T_+$ — in fact, that example will have a one-parameter family of distinct such maps. On the other hand, the arguments in Section 5.4 will show that a Λ_- isometry is determined by where a single point of \tilde{G} goes in T_+ .

Proof of Lemma 2.20. The first part follows by combining Corollary 2.14 and Lemma 2.17.

To prove uniqueness of the realization in T_+ of a leaf $\tilde{\ell}$ of $\tilde{\Lambda}_-$, we must prove that $f_g(\tilde{\ell})$ is well-defined independent of g . Think of $\tilde{\ell}$ abstractly as an element of \mathcal{GF}_r , and choose a sequence $\gamma_i \in F_r$ whose abstract axes $A(\gamma_i)$ in \mathcal{GF}_r limit to $\tilde{\ell}$. For any train track representative $g: G \rightarrow G$ of ϕ , the realizations $A(\gamma_i)_{\tilde{G}}$ converge to the realization $\tilde{\ell}_G$ in the Gromov-Hausdorff topology on closed subsets of \tilde{G} . It follows that the sets $p_i = \tilde{\ell}_G \cap A(\gamma_i)_{\tilde{G}}$ form a longer and longer family of paths that exhausts $\tilde{\ell}_G$. While $A(\gamma_i)_{\tilde{G}}$ need not be legal

in \tilde{G} and so may not embed in T_+ under f_g , its subpath p_i is legal and does embed in T_+ . By bounded cancellation, Lemma 2.5, there exists B independent of i and depending only on f_g , such that if p'_i is obtained from p_i by cutting off segments of length B at both ends, then $f_g(p'_i)$ embeds in $A(\gamma_i)_{T_+}$, the axis of γ_i in T_+ . But the sequence of paths $f_g(p'_i)$ form a longer and longer family of paths in T_+ that exhausts $f_g(\tilde{\ell}_G)$. It follows that $A(\gamma_i)_{T_+}$ approaches $f_g(\tilde{\ell}_G)$ in the Gromov-Hausdorff topology on closed subsets of T_+ . Since the sequence γ_i is independent of the choice of $g: G \rightarrow G$, uniqueness of $f_g(\tilde{\ell})$ in T_+ follows. \diamond

We conclude with a relation between Λ_- and the homothety $\Phi_+: T_+ \rightarrow T_+$.

Corollary 2.21. *For each $\Phi \in \text{Aut}(F_n)$ representing ϕ , and each leaf $\tilde{\ell}$ of $\tilde{\Lambda}_-$, the map Φ_+ takes the realization of $\tilde{\ell}$ in T_+ to the realization of $\Phi_{\#}(\tilde{\ell})$ in T_+ .*

Proof. Let $g: G \rightarrow G$ be a train track representative of ϕ and $\check{g}: \tilde{G} \rightarrow \tilde{G}$ the lift corresponding to Φ . Let $\tilde{\ell}$ denote its own realization in \tilde{G} , and so $f_g(\tilde{\ell})$ is its realization in T_+ . Also, the realization of $\Phi_{\#}(\tilde{\ell})$ in \tilde{G} is $\check{g}(\tilde{\ell})$, and its realization in T_+ is $f_g(\check{g}(\tilde{\ell}))$. But we know that

$$f_g \circ \check{g} = \Phi_+ \circ f_g$$

and so Φ_+ takes the realization of $\tilde{\ell}$ in T_+ to the realization of $\Phi_{\#}(\tilde{\ell})$ in T_+ . \diamond

3 The ideal Whitehead graph

On a closed oriented surface S , the topological structure of the unstable foliation of a pseudo-Anosov homeomorphism $\phi: S \rightarrow S$ near a singular point $s \in S$ is fully determined by the number $n \geq 3$ of unstable prongs at s . Associated to s there is a principal region (that is, a complementary component) of the unstable geodesic lamination of ϕ , this region is an ideal n -gon, and its n sides are the singular leaves associated to s , which are circularly ordered by the asymptote relation. Also, in the \mathbf{R} -tree dual to the unstable foliation, associated to each n -pronged singularity s there is an orbit of branch points of the \mathbf{R} -tree, and the number of directions at each such branch point is exactly n . The singularity structure of ϕ is therefore described fully by a list of integers n_i giving the number of unstable prongs at each singularity. The singularity structure is constrained by the Euler-Poincaré index equation $\sum(n_i - 2) = -2\chi(S)$, and the paper [MS93] proves that outside of a few fully understood exceptions that occur only for a few surfaces of low topological complexity, the index equation gives a complete description of all singularity structures that actually occur for pseudo-Anosov homeomorphisms of S .

Singularity structures of nongeometric, fully irreducible outer automorphisms $\phi \in F_r$ are less well understood. The most that is known so far is an index inequality of Euler-Poincaré type that is proved in [GJLL98]: there are finitely many F_r orbits of branch points in T_+^ϕ , each branch point has finitely many directions, and if n_i is the number of directions at each branch point in the i^{th} orbit then $\sum(n_i - 2) \leq 2(r - 1)$. Whether this gives a complete description of all singularity structures that actually occur, along the lines of [MS93], has not been investigated. In the index inequality of [GJLL98], the significant information about a singularity is, like in the surface case, a raw count of its number of “unstable directions”, encoded as directions in T_+^ϕ .

In this section we study a finer invariant of the singularity structure of a nongeometric, fully irreducible $\phi \in \text{Out}(F_r)$, the “ideal Whitehead graph” of ϕ , which captures the asymptotic relations among singular leaves of $\Lambda_-(\phi)$. The ideal Whitehead graph is first defined in terms of principal automorphisms representing ϕ or, equivalently, in terms of principal lifts of a train track representative. The main results of this section give alternative descriptions of the ideal Whitehead graph: in terms of singular leaves of train track representatives (Lemma 3.1 and Corollary 3.2); next in terms of the asymptote relation among singular leaves (Lemma 3.3); and finally in terms of branch points in the \mathbf{R} -tree T_+^ϕ (Lemma 3.4).

To describe the ideal Whitehead graph very roughly, to each “singularity of ϕ ” there is associated a component of the ideal Whitehead graph, a finite graph with singular leaves as edges, and with vertices describing asymptotic relations between singular leaves. By analogy, for a pseudo-Anosov surface singularity with n unstable prongs, the corresponding component of the ideal Whitehead graph has n vertices and n edges forming a topological circle. For a general nongeometric, fully irreducible $\phi \in \text{Out}(F_r)$, components of the ideal Whitehead graph need not be topological circles; indeed, they may have cut points, as we show by example in Section 3.4.

For present purposes we restrict to the case that ϕ is rotationless, but the reader will easily see how to proceed without this restriction. Also, we shall restrict attention to an F_r -equivariant version of the ideal Whitehead graph; each component having trivial stabilizer, the reader will easily see how to pass to a quotient under F_r to obtain a nonequivariant version.

For the rest of this section we fix a nongeometric, fully irreducible, rotationless $\phi \in \text{Out}(F_r)$, and we denote $\Lambda_- = \Lambda_-(\phi)$ and $T_+ = T_+^\phi$. All automorphisms and all train track maps will be representatives of ϕ .

3.1 Definition and structure of the ideal Whitehead graph

Recall that, abstractly, leaves of Λ_- are just points $(P_1, P_2) \in D^2 \partial F_r$.

For every principal automorphism Φ let $L(\Phi)$ be the set of leaves (P_1, P_2) of Λ_- with $P_1, P_2 \in \text{Fix}_N(\hat{\Phi})$. Define the *component* $W(\Phi)$ of the ideal Whitehead graph determined by Φ to be the graph with one vertex for each point in $\text{Fix}_N(\hat{\Phi})$, and an edge connecting $P_1, P_2 \in \text{Fix}_N(\hat{\Phi})$ if there is a leaf $(P_1, P_2) \in L(\Phi)$. Corollary 2.9 implies that for each point $P \in \partial F_r$ there is at most one principle automorphism Φ such that $P \in \text{Fix}_N(\hat{\Phi})$, and so $L(\Phi_1) \cap L(\Phi_2) = \emptyset$ for distinct principle automorphisms Φ_1, Φ_2 . It therefore makes sense to define the *ideal Whitehead graph* $W(\phi)$ to be the disjoint union of its components $W(\Phi)$, one for each principal automorphism Φ representing ϕ .

The definition of $W(\phi)$ can be easily reformulated in terms of any train track representative $g: G \rightarrow G$ of ϕ . For every principal lift $\check{g}: \tilde{G} \rightarrow \tilde{G}$, let $L(\check{g})$ denote the set of leaves of $\tilde{\Lambda}_-$ with both ideal endpoints in $\text{Fix}_N(\hat{g})$, define $W(\check{g})$ to be the graph with one vertex for each point in $\text{Fix}_N(\hat{g})$ and an edge connecting the vertices corresponding to $P_1, P_2 \in \text{Fix}_N(\hat{g})$ if there is a leaf $\ell \in L(\check{g})$ with ideal endpoints P_1 and P_2 , and define $W(g)$ to be the disjoint union of the $W(\check{g})$. Note that if Φ is the principal automorphism corresponding to \check{g} then there are natural isomorphisms $L(\check{g}) \approx L(\Phi)$ and $W(\check{g}) \approx W(\Phi)$, and there is a natural isomorphism $W(g) \approx W(\phi)$. Formally, it is convenient to let the vertex set of $W(\check{g})$ be $\text{Fix}_N(\hat{g})$ and to let the edge interiors of $W(\check{g})$ be the disjoint union of the leaves of $L(\check{g})$.

For any principal lift $\check{g}: \tilde{G} \rightarrow \tilde{G}$ and $\tilde{v} \in \text{Fix}(\check{g})$, we identify the local stable Whitehead graph $SW_{\tilde{v}}$ with a subgraph of $W(\check{g})$ as follows. Each vertex d_i of $SW_{\tilde{v}}$ is a direction at \tilde{v} fixed by $D\check{g}$, d_i is the initial direction of a unique singular ray R_i , and d_i is identified with the vertex of $W(\check{g})$ which is the ideal endpoint P_i of the ray R_i . As shown in Section 2.7, two vertices d_i, d_j of $SW_{\tilde{v}}$ that are connected by an edge of $SW_{\tilde{v}}$ correspond to a singular leaf $\ell = R_1 \cup R_2$ with ideal endpoints P_1, P_2 , and so ℓ corresponds to an edge of $W(\check{g})$ with endpoints P_1, P_2 .

Lemma 3.1. *Let \check{g} be a principal lift of g .*

- (1) *Every $P \in \text{Fix}_N(\hat{g})$ is the endpoint of at least one element of $L(\check{g})$.*
- (2) *Singular rays R_1 at $\tilde{v}_1 \in \text{Fix}(\check{g})$ and R_2 at $\tilde{v}_2 \in \text{Fix}(\check{g})$ terminate at the same point $P \in \text{Fix}(\hat{g})$ if and only if the path $\tilde{\sigma}$ connecting \tilde{v}_1 to \tilde{v}_2 projects to an indivisible Nielsen path σ and determines the same direction at \tilde{v}_i as does R_i .*
- (3) *$L(\check{g})$ is the disjoint union, over all $\tilde{v} \in \text{Fix}(\check{g})$, of the singular leaves at \tilde{v} .*

Proof. Item (1) follows from Lemma 4.7 of [FH04].

The path σ of (2) is a Nielsen path. If R_1 and R_2 eventually coincide then $\tilde{\sigma} = \tilde{\alpha}\tilde{\beta}^{-1}$ where $\tilde{\alpha} \subset R_1$ and $\tilde{\beta} \subset R_2$. This proves the only if part of (2). For the if part, suppose that $\sigma = \alpha\bar{\beta}$ is indivisible. Let $\tilde{\alpha}$ and $\tilde{\beta}$ be the lifts of α and β that begin at \tilde{v}_1 and \tilde{v}_2 respectively. For all $n \geq 1$, there are paths $\tilde{\tau}_n$ such that $\check{g}^n(\tilde{\alpha}) = \tilde{\alpha}\tilde{\tau}_n$ and $\check{g}^n(\tilde{\beta}) = \tilde{\beta}\tilde{\tau}_n$.

The paths $\check{g}^n(\tilde{\alpha})$ are an increasing sequence whose union is R_1 and the paths $\check{g}^n(\tilde{\beta})$ are an increasing sequence whose union is R_2 . This completes the proof of (2).

Each singular leaf based at an element of $\text{Fix}(\check{g})$ is contained in $L(\check{g})$ by construction. Conversely, since the endpoints of any $\tilde{\ell} \in L(\check{g})$ are attractors for \hat{g} , there is a necessarily unique fixed point $\tilde{x} \in \tilde{l}$; see for example Lemma 3.9 and Lemma 3.10 of [FH04]. Lemma 2.11 implies that \tilde{x} is a principal vertex. The turn taken by $\tilde{\ell}$ at \tilde{x} determines an edge of $SW_{\tilde{x}}$ and hence a singular leaf $\tilde{\ell}'$ at \tilde{x} . If $\tilde{\ell} \neq \tilde{\ell}'$ then there are rays R and R' based at \tilde{x} with fixed point free interiors, with a common initial edge and with distinct terminal endpoints P and P' in $\text{Fix}_N(\hat{g})$. The last point of $R \cap R'$ is therefore mapped by \check{g} both to $R - R'$ and to $R' - R$, a contradiction. This proves that ℓ is a singular leaf at \tilde{x} and completes the proof of (3). \diamond

The following corollary says that each component of the ideal Whitehead graph is a union of local Whitehead graphs, pasted together along cut points:

Corollary 3.2. *Let \check{g} be a principal lift of g .*

- (1) $W(\check{g})$ is connected.
- (2) $W(\check{g}) = \cup SW(\tilde{v})$ where the union is taken over all $\tilde{v} \in \text{Fix}(\check{g})$.
- (3) For $i \neq j$, $SW(\tilde{v}_i)$ and $SW(\tilde{v}_j)$ intersect in at most one vertex. If they do intersect at a vertex P , then P is a cut point of $W(\check{g})$, in fact P separates $SW(\tilde{v}_i)$ from $SW(\tilde{v}_j)$ in $W(\check{g})$.

Proof. (2) follows from the first and third item of Lemma 3.1. (1) follows from (2), item 2 of Lemma 3.1, and the fact that any two elements of $\text{Fix}(\check{g})$ are connected by a finite concatenation of indivisible Nielsen paths.

To simplify notation, for proving (3) we assume $i = 1, j = 2$. Suppose that $P \in \text{Fix}_N(\hat{g})$ is a common vertex of $SW(\tilde{v}_1) \cap SW(\tilde{v}_2)$. There are singular rays R_1 and R_2 terminating at P and initiating at \tilde{v}_1 and \tilde{v}_2 respectively. Since no leaf of $\tilde{\Lambda}_-$ contains both \tilde{v}_1 and \tilde{v}_2 , it follows that $R_1 \cap R_2 \cap \{\tilde{v}_1, \tilde{v}_2\} = \emptyset$. If $Q \neq P$ is another common vertex of $SW(\tilde{v}_1) \cap SW(\tilde{v}_2)$ then there are singular rays R'_1 and R'_2 terminating at Q and initiating at \tilde{v}_1 and \tilde{v}_2 respectively. The initial direction of R_1 is distinct from that of R'_1 and similarly for R_2 and R'_2 . It follows that $R_1 \cup R_2 \cup R'_1 \cup R'_2$ contains an embedded circle, contradicting that \tilde{G} is a tree. This proves $SW(\tilde{v}_1) \cap SW(\tilde{v}_2) = \{P\}$.

To prepare for proving the last part of (3), let $W^d(\check{g})$ denote $W(\check{g})$ with vertices deleted, namely, the disjoint union of all singular leaves based at all fixed points of \check{g} ; formally, one can think of $W^d(\check{g})$ as the set of ordered pairs (l, x) where l is a singular leaf based at a fixed point of \check{g} and $x \in l$. Let $r: W^d(\check{g}) \rightarrow \tilde{G}$ be the natural map, taking each singular leaf to itself by the identity map; formally, $r(l, x) = x$. The set $r^{-1}(\text{Fix}(\check{g}))$ consists of exactly one point in each edge of $W(\check{g})$, called the *basepoint* of that edge, separating the edge into two *open half-edges*. Suppose that P does not separate $SW(\tilde{v}_1)$ from $SW(\tilde{v}_2)$ in $W(\check{g})$. Then there are edge basepoints $w_1, w_2 \in W^d(\check{g})$ with $r(w_i) = v_i$, and there is an embedded path γ in $W(\check{g})$ from w_1 to w_2 disjoint from P . Let $\gamma^d = \gamma \cap W^d(\check{g})$, that is, γ with vertices of $W(\check{g})$ deleted. The image $r(\gamma^d)$ is a connected subset of \tilde{G} , because near any deleted vertex Q of γ the two ends of γ incident to Q are mapped by r to rays in \tilde{G} asymptotic to Q . Also, since

γ does not pass through P , and since γ is embedded, it follows that γ does not intersect any open half-edge incident to P , from which it follows that the image $r(\gamma^d)$ is disjoint from $\text{int}(R_1) \cup \text{int}(R_2)$. Therefore, $r(\gamma^d)$ contains an embedded path in $\tilde{G} - (\text{int}(R_1) \cup \text{int}(R_2))$ from \tilde{v}_1 to \tilde{v}_2 . It follows that $r(\gamma^d) \cup R_1 \cup R_2$ contains an embedded circle, contradicting that \tilde{G} is a tree. \diamond

3.2 Asymptotic leaves and the ideal Whitehead graph

A pair of distinct leaves of $\tilde{\Lambda}_-$ are *asymptotic* if they have a common infinite end. This generates an equivalence relation on leaves of $\tilde{\Lambda}_-$ whose non-trivial classes are called *asymptotic equivalence classes* of $\tilde{\Lambda}_-$.

Our next lemma relates a component of the ideal Whitehead graph to asymptotic equivalence classes.

Lemma 3.3. *The map $\check{g} \mapsto L(\check{g})$ defines a bijection between the set of principal lifts of g and the set of asymptotic equivalence classes of $\tilde{\Lambda}_-$. Two ends of edges of $W(\check{g})$ are incident to the same vertex if and only if the corresponding rays in leaves of $L(\check{g})$ are asymptotic.*

Proof. The second sentence is obvious.

Injectivity of the map $\check{g} \mapsto L(\check{g})$ follows from the fact that $L(\check{g}_1) \cap L(\check{g}_2) = \emptyset$ for distinct \check{g}_1 and \check{g}_2 . Singular leaves at $\tilde{v} \in \text{Fix}(\check{g})$ belong to the same asymptotic equivalence class because $SW_{\tilde{v}}$ is connected. Since any two elements of $\text{Fix}(\check{g})$ are connected by a finite concatenation of indivisible Nielsen paths, the second item of Lemma 3.1 implies that all elements of $L(\check{g})$ are in the same asymptotic equivalence class.

It remains to show that if $\tilde{\ell}'$ and $\tilde{\ell}''$ are asymptotic leaves in \tilde{G} then they are elements of some $L(\check{g})$. It is convenient to work in G rather than \tilde{G} so consider asymptotic leaves $\ell', \ell'' \in \Lambda_-$. For each $j \geq 0$, there are leaves ℓ'_j and ℓ''_j such that $g^j(\ell'_j) = \ell'$ and $g^j(\ell''_j) = \ell''$. The intersection of ℓ'_j and ℓ''_j is a ray R_j based at a vertex x_j . Denote the rays in ℓ'_j and ℓ''_j that are complementary to R_j by R'_j and R''_j respectively. Let E_j, E'_j and E''_j be the initial edges of R_j, R'_j and R''_j respectively and let d_j, d'_j and d''_j be the directions that they determine.

As a first case suppose that (d'_j, d''_j) is legal for all $j \geq 0$. Then x_j is principal and $g^i(x_{j+i}) = x_j$ for all $i, j \geq 0$. It follows that $x_j \in \text{Fix}(g)$ is independent of j and we now denote x_j by v . The analogous argument shows that E_j, E'_j and E''_j are independent of j and we denote these by E, E' and E'' . For all j , the paths $g^j(E)$ and $g^j(E')$ are respectively initial subpaths of R_0 and R'_0 . It follows that ℓ' is the singular leaf at v determined by E and E' . The symmetric argument shows that ℓ'' is the singular leaf at v determined by E and E'' . The point \tilde{v} at which $\tilde{\ell}'$ and $\tilde{\ell}''$ diverge projects to v . If \check{g} is the lift that fixes \tilde{v} , then ℓ' and ℓ'' are singular leaves at \tilde{v} for \check{g} . This completes the proof in this first case.

Assume now that (d'_j, d''_j) is illegal for some, and hence every sufficiently large, j . Item 5 of Lemma 2.7 and Lemma 2.12 imply that there is an indivisible Nielsen path $\alpha_j \bar{\beta}_j$ such that $\bar{\alpha}_j$ is an initial path of R'_j , $\bar{\beta}_j$ is an initial path of R''_j , and $g_{\#}^i(\alpha_{i+j} \bar{\beta}_{i+j}) = \alpha_j \bar{\beta}_j$ for all $i, j \geq 0$. It follows that α_j, β_j, x_j and the initial endpoints y_j and z_j of α_j and β_j respectively are independent of j ; denote these by α, β, v, y and z . Arguing as in the previous case we see that ℓ' and ℓ'' are singular leaves at y and z respectively. The point \tilde{v} at which $\tilde{\ell}'$ and $\tilde{\ell}''$ diverge projects to the common terminal endpoint v of α and β . Let $\tilde{\alpha}$ and $\tilde{\beta}$ be the lifts

that terminate at \tilde{v} and let \check{g} be the lift that fixes the initial endpoints \tilde{y} and \tilde{z} of $\tilde{\alpha}$ and $\tilde{\beta}$. Then ℓ' and ℓ'' are singular leaves for \check{g} at \tilde{y} and \tilde{z} respectively. \diamond

3.3 T_+ and the ideal Whitehead graph

Next we relate the ideal Whitehead graph to T_+ . Recall Corollary 2.20 which explains how the leaves of $\tilde{\Lambda}_-$ are realized in T_+ .

Suppose that $b \in T_+$ is a branch point, $\Phi \in \text{PA}(\phi)$ is the principal automorphism corresponding to b by Lemma 2.16, and $\check{g}: \tilde{G} \rightarrow \tilde{G}$ is the corresponding principal lift. Each singular leaf ℓ in $L(\check{g})$ passes through some $x \in \text{Fix}(\check{g})$, and so $f_g(\ell)$, the realization of ℓ in T_+ , passes through the point $f_g(x) = b$; we call these the *singular leaves at b* . Each singular leaf at b is divided by b into *singular rays at b* which are the images under f_g of the singular rays at some $x \in \text{Fix}(\check{g})$.

Define the *local Whitehead graph* at a point $z \in T_+$, denoted $W(z; T_+)$, to be the graph with one vertex for each direction d_i at z and with an edge connecting d_i to d_j if the turn (d_i, d_j) is taken by the realization of some leaf of Λ_- in T_+ . For each $x \in \tilde{G}$ the map $D_x f_g$ maps the directions of \tilde{G} at x to the directions of T_+ at $f(b)$, and this map extends to a simplicial map $W(x; \tilde{G}) \mapsto W(b; T_+)$, still denoted $D_x f_g$ (though we often drop x from the notation $D_x f_g$).

With b, Φ, \check{g} as above, for each $x \in \text{Fix}(\check{g})$ the map $D_x f_g$ is injective on the set of fixed directions at x , and therefore $D_x f_g$ restricts to a simplicial embedding $SW(x; \tilde{G}) \hookrightarrow W(b; T_+)$.

Lemma 3.4. *With b, Φ , and \check{g} as above, as x varies over $\text{Fix}(\check{g})$ the embeddings*

$$D_x f_g: SW(x; \tilde{G}) \hookrightarrow W(b; T_+)$$

induce an isomorphism of $W(\Phi)$ with $W(b; T_+)$. More precisely,

- (1) *Every direction at b is realized by a singular ray.*
- (2) *Singular rays R_1 and R_2 in \tilde{G} with basepoints in $\text{Fix}(\check{g})$ map to the same singular ray in T_+ if and only if they terminate at the same point in $\text{Fix}_N(\hat{g})$.*
- (3) *A turn at b is taken by some leaf of $\tilde{\Lambda}_-$ if and only if it is taken by a singular leaf at b .*

Proof. Observe first that, by Corollary 3.2 and (2), the maps $D_x f_g: SW(x; \tilde{G}) \rightarrow W(b; T_+)$ for $x \in \text{Fix}(\check{g})$ fit together to define a simplicial map $W(\Phi) \mapsto W(b; T_+)$, which by items (1)–(3) is an isomorphism. Item (1) follows from the second item of Lemma 2.16 and the fact that every fixed direction at an element of $\text{Fix}(\check{g})$ is realized by a singular ray.

The “if” part of (2) follows from the second item of Lemma 3.1 and Corollary 2.14. For the “only if” part, let R_1 and R_2 be singular rays at \tilde{v}_1 and \tilde{v}_2 such that $f_g(R_1) = f_g(R_2)$. We prove that R_1 and R_2 converge to the same point in $\text{Fix}(\hat{g})$. Let d_i be the direction determined by R_i at \tilde{v}_i . If $\tilde{v}_1 = \tilde{v}_2$, then $R_1 = R_2$ and the result is obvious. Suppose then that $\tilde{v}_1 \neq \tilde{v}_2$. The path connecting \tilde{v}_1 to \tilde{v}_2 splits as a concatenation $\tilde{\mu} = \tilde{\mu}_1 \cdot \dots \cdot \tilde{\mu}_l$ where each μ_i is an indivisible Nielsen path. For $i = 1, 2$ choose \tilde{x}_i close to and in the direction d_i from \tilde{v}_i such that $f_g(\tilde{x}_1) = f_g(\tilde{x}_2)$. The path $\tilde{\mu}'$ connecting \tilde{x}_1 to \tilde{x}_2 is obtained from $\tilde{\mu}$ by

making alterations at the beginning and end. If d_1 is the direction determined by $\tilde{\mu}_1$, then the path from \tilde{v}_1 to \tilde{x}_1 is removed from $\tilde{\mu}_1$; otherwise the path from \tilde{x}_1 to \tilde{v}_1 is added. Similar changes are made at the terminal end. Since there are only finitely many indivisible Nielsen or pre-Nielsen paths in G we may assume that \tilde{x}_1 and \tilde{x}_2 have been chosen so that $\tilde{\mu}'$ does not project to a Nielsen or pre-Nielsen path. Corollary 2.14 implies that $\check{g}^k(\tilde{x}_1) = \check{g}^k(\tilde{x}_2)$ for some $k > 0$. The second item of Lemma 2.7 therefore implies that \tilde{x}_1 and \tilde{x}_2 belong to the interior of the same $\tilde{\mu}_j$. It follows that $\mu = \mu_1$ is indivisible and that d_1 and d_2 point into $\tilde{\mu}$. The second item of Lemma 3.1 implies that R_1 and R_2 terminate at the same endpoint as desired.

The “if” part of (3) is obvious. Suppose the turn (d_1^+, d_2^+) is taken by a leaf $\tilde{\ell}$ of $\tilde{\Lambda}_-$. The realization of $\tilde{\ell}$ in \tilde{G} contains some $\tilde{x} \in f_g^{-1}(b)$. By Lemma 2.16(3), for some $m > 0$ we have $\check{g}^m(\tilde{x}) \in \text{Fix}(\check{g})$ and the two directions d_1, d_2 of $\check{g}^m(\tilde{\ell})$ at $\check{g}^m(\tilde{x})$ are fixed by $D\check{g}$, so $f_g(g^m(\tilde{x})) = b$ and $Df_g(d_i) = d_i^+$. Also, since the turn (d_1, d_2) is taken by some leaf, namely $\check{g}^m(\tilde{\ell})$, it follows that this turn is taken by some singular leaf $\tilde{\ell}'$ based at $\check{g}^m(\tilde{x})$, and so $f_g(\ell')$, the realization of $\tilde{\ell}'$ in T_+ , is a singular leaf based at b taking the turn (d_1^+, d_2^+) . \diamond

3.4 An example of an ideal Whitehead graph

In this section we give an example of a train track map representing a nongeometric, fully irreducible element $\phi \in \text{Out}(F_3)$, together with a complete computation of (the quotient by F_3 of) the ideal Whitehead graph of ϕ .

Example: Consider $\phi \in \text{Out}(F_3)$ defined by the automorphism

$$\begin{aligned} A &\mapsto AFGFGF \\ \Phi : F &\mapsto FGF \\ G &\mapsto GFAGF \end{aligned}$$


That Φ is indeed an automorphism is easily checked by computing its inverse. Direct calculation shows that the eigenvalues of the matrix of the abelianization of Φ , accurate to the closest 2nd decimal place, are 4.08 and $.46 \pm .18i$. No eigenvalue is a root of unity and so the following lemma, whose proof is found below, applies to each power of ϕ to show that ϕ is fully irreducible.

Lemma 3.5. *If $\phi \in \text{Out}(F_3)$ is reducible then at least one eigenvalue of the abelianization of ϕ is a root of unity.*

The formula for Φ also defines an irreducible train track representative g of ϕ acting on the rose with edges labelled A, F, G . We will see shortly that g has no periodic Nielsen path. It follows by [BH92] that ϕ is not geometric. It also follows, by Lemma 3.1, that the local stable Whitehead graph at the valence 6 vertex of the rose is equal to the ideal Whitehead graph of ϕ (modulo the action of F_3).

The first of the following two graphs is the local Whitehead graph of g at the valence 6 vertex. The second is the local stable Whitehead graph, as well as the ideal Whitehead graph

of ϕ (modulo F_3), as we shall soon show:



Note that the edge $\{\overline{F}, F\}$ only appears in the local Whitehead graph after the second iterate of g , and that \overline{A} is the only nonperiodic direction.

To show that the second graph above is the ideal Whitehead graph of ϕ (modulo F_3), it remains to demonstrate that g has no periodic indivisible Nielsen path $\alpha * \beta$. This demonstration amounts to a recursive construction of longer and longer candidates for terminal segments of α and β , throwing out impossible cases whenever encountered, until either α, β are constructed, or all cases are exhausted and $\alpha * \beta$ is thereby proved not to exist.

Suppose that there exists an indivisible periodic Nielsen path $\alpha * \beta$. Since $\{\overline{A}, \overline{F}\}$ is the unique illegal turn of g , up to relabelling we have

$$\begin{aligned}
 (1) \quad & \alpha = \alpha_1 * A & g\alpha &= g\alpha_1 * AFGFGF \\
 & \beta = \beta_1 * F & g\beta &= g\beta_1 * FGF
 \end{aligned}$$

The last letter of $g\beta_1$ must be G , because the last letters of $g\alpha, g\beta$ not in their common terminal segment must be A, F . The last letter of β_1 must therefore be G :

$$\begin{aligned}
 (2) \quad & \alpha = \alpha_2 * A & g\alpha &= g\alpha_2 * AFGFGF \\
 & \beta = \beta_2 * GF & g\beta &= g\beta_2 * GFAFGFGF
 \end{aligned}$$

The recursion continues, with justifications to follow:

$$\begin{array}{lcl}
 (3) & \alpha = \alpha_3 * FA & g\alpha = g\alpha_3 * FGFAFGFGF \\
 & \beta = \beta_3 * GF & g\beta = g\beta_3 * GFAFGFGF \\
 \hline
 (4) & \alpha = \alpha_4 * FA & g\alpha = g\alpha_4 * FGFAFGFGF \\
 & \beta = \beta_4 * FGF & g\beta = g\beta_4 * FGFGFAFGFGF \\
 \hline
 (5) & \alpha = \alpha_5 * GFA & g\alpha = g\alpha_5 * GFAFGFGFAFGFGF \\
 & \beta = \beta_5 * FGF & g\beta = g\beta_5 * FGFGFAFGFGF
 \end{array}$$

To justify (3), $g\alpha_2$ must end in A or F , but no edge has image under g ending in A and so $g\alpha_2$ ends in F ; this implies that α_2 ends in A or F , but AA is not a leaf segment (because $\{\overline{A}, A\}$ is not an edge of the local Whitehead graph) and so α_2 ends in F . The justification of (4) is similar to (3), and (5) is similar to (2).

Next we conclude as in (3) that $g\beta_5$ ends in F , and so we branch into two cases: (5.1) β_5 ends in A ; or (5.2) β_5 ends in F .

$$\begin{array}{lcl}
 (5.1) & \alpha = \alpha_5 * GFA & g\alpha = g\alpha_5 * GFA.FGFGFAFGFGF \\
 & \beta = \beta_{5.1} * AFGF & g\beta = g\beta_{5.1} * AFGFGF.FGFGFAFGFGF \\
 \hline
 (5.2) & \alpha = \alpha_5 * GFA & g\alpha = g\alpha_5 * GFA.FGFGFAFGFGF \\
 & \beta = \beta_{5.2} * FFGF & g\beta = g\beta_{5.2} * FGF.FGFGFAFGFGF
 \end{array}$$

In these examples the dot indicates the beginning of the longest common terminal segment. If $\alpha * \beta$ were a fixed Nielsen path then it would follow that $\alpha_5 * GFA = g\alpha_5 * GFA$ implying $\alpha_5 = g\alpha_5$ and we conclude that α_5 is the empty word. But the same conclusion holds even under our presumption that $\alpha * \beta$ is a periodic Nielsen path of some period $k \geq 1$, because repeating the analysis above allows us to conclude that for all $i \geq 1$, the edge path $g^i \alpha_5 * GFA$ is the initial segment of $g^i \alpha$ obtained by truncating the longest common terminal segment of $g^i \alpha$ and $g^i \beta$, and so $\alpha_5 = g^k \alpha_5$ implying that α_5 is the empty word.

We have shown that $\alpha = GFA$ is one-half of any periodic indivisible Nielsen path. Next we show that any periodic indivisible Nielsen path has period $k = 1$. For suppose that $k > 1$ were the period of $\alpha * \beta_1$. Straightening $g(\alpha * \beta_1)$ would yield $\alpha * \beta_2$ with β_1, β_2 distinct from each other and from α . This implies that the path obtained by straightening $\beta_1 * \beta_2$ is a periodic indivisible Nielsen path, and so $\beta_1 = \gamma_1 * \delta$, $\beta_2 = \gamma_2 * \delta$ with $\gamma_1 * \gamma_2$ a periodic indivisible Nielsen path. But it follows that $\gamma_i = \alpha = GFA$ for some $i = 1, 2$, which implies that $\text{Len}(GFA) = \text{Len}(\alpha) = \text{Len}(\gamma_i) \leq \text{Len}(\beta_i) = \text{Len}(\alpha)$ and so the lengths are equal and $\alpha = \gamma_i = \beta_i$, contradicting that $\alpha \neq \beta_i$.

We now continue analyzing cases (5.1) and (5.2), under the stronger condition that $\alpha * \beta$ is a *fixed* indivisible Nielsen path. The segment of $g\beta$ before the dot must agree with β , as it does in (5.2) but not in (5.1), and so (5.1) is ruled out. Continuing with (5.2), since $g\beta$ before the dot must agree with β , the last letter of $g\beta_{5.2}$ must be F , and so either (5.1.1) the last letter of $\beta_{5.2}$ is A , or (5.1.2) the last letter of $\beta_{5.2}$ is F :

$$\begin{array}{lcl} (5.2.1) & \alpha = GFA & g\alpha = GFA.FGFGFAFGFGF \\ & \beta = \beta_{5.2.1} * AFFGF & g\beta = g\beta_{5.2.1} * AFGFGFFGF.FGFGFAFGFGF \\ \hline (5.2.2) & \alpha = GFA & g\alpha = GFA.FGFGFAFGFGF \\ & \beta = \beta_{5.2.2} * FFFGF & g\beta = g\beta_{5.2.2} * FGFFGF.FGFGFAFGFGF \end{array}$$

In neither (5.2.1) nor (5.2.2) does β agree with $g\beta$ before the dot, and so both cases are ruled out. All cases being exhausted, it follows that g has no periodic indivisible Nielsen path.

To complete the example we give:

Proof of Lemma 3.5. Since ϕ is reducible there is a free factor F of rank one or of rank two whose conjugacy class is invariant under an iterate of ϕ . After passing to a power of ϕ , the conjugacy class of F becomes ϕ -invariant. If F has rank one then the conjugacy class of a generator of F is ϕ -invariant, and the image of this generator in \mathbf{Z}^3 is an eigenvector with eigenvalue one. We may therefore assume that F has rank two. Lemma 3.2 of [BFH00] implies that, after passing to a power of ϕ , there exist $\Phi \in \text{Aut}(F_3)$ representing ϕ and generators a, b and c of F_3 such that a and b generate F and such that $\Phi(c) = ucv$ for some $u, v \in F$. Letting x be the abelianized image of c and Z^2 the abelianized image of F , we have $\phi_\#(x) = x + s$ for some $s \in \mathbf{Z}^2$. If 1 is an eigenvalue for the restriction of $\phi_\#$ to \mathbf{Z}^2 , we are done. Otherwise there exists $t \in \mathbf{Z}^2$ such that $\phi_\#(t) = t - s$ and $x + t$ is an eigenvector with eigenvalue one. \diamond

4 Cutting and pasting local stable Whitehead graphs

Let $\phi \in \text{Out}(F_n)$ be nongeometric and fully irreducible. Corollary 3.2 shows that choosing a train track representative $g: G \rightarrow G$ of ϕ induces a decomposition of each component of the ideal Whitehead graph $W(\phi)$, expressing that component as the union of local stable Whitehead graphs pasted together along cut vertices. We refer to this as the *local decomposition* of $W(\phi)$ induced by g .

In this section we study how the local decomposition changes under certain moves that alter the train track representative. Our main result, Corollary 4.4, explains how to construct a train track representative so that the local decomposition is as fine as possible, having no cut vertices in any local stable Whitehead graph. To put it another way, every finite graph has a unique decomposition into a union of maximal subgraphs without cut vertices, called the *cut point decomposition* of the graph, and Corollary 4.4 produces a train track representative that simultaneously realizes the cut point decomposition for all components of $W(\phi)$.

The local decomposition of $W(\phi)$ is intimately connected with Nielsen classes of principal vertices of train track representatives. Given a train track representative $g: G \rightarrow G$ of ϕ , while the number of Nielsen classes of principal vertices of g depends only on ϕ , the number of principal vertices depends on the choice of $g: G \rightarrow G$. As Lemma 3.1 and Corollary 3.2 show, for any train track representative $g: G \rightarrow G$ of ϕ , the decomposition of $W(\phi)$ into local stable Whitehead graphs of g pasted together along cut vertices corresponds to the decomposition of Nielsen classes of principal vertices of g into individual vertices joined by Nielsen paths.

First, in Lemma 4.1, we describe a condition which allows us to reduce the number of principal vertices, and a method for doing so — this corresponds to coarsening the local decomposition of $W(\phi)$. Then, in Lemma 4.3, we give necessary and sufficient conditions for recognizing when a given irreducible train track map is the product of this method. In this case the method can be reversed to increase the number of principal vertices, refining the local decomposition of $W(\phi)$. As an application, in Corollary 4.4 we show that the finest decomposition of $W(\phi)$ is realized by maximizing the number of principal vertices of a train track representative; the maximum exists because the number of principal vertices is bounded above by the number of singular leaves of Λ_- .

4.1 Pasting local stable Whitehead graphs

Close study of Lemma 3.1 and Corollary 3.2 indicates that we should be able to paste together local stable Whitehead graphs and coarsen the local decomposition of the ideal Whitehead graph, by decreasing the number of principal vertices, or equivalently by eliminating Nielsen paths. The following describes a process for carrying out this elimination; it is a variation on Lemma 3.8 of [BH92].

Lemma 4.1. *Suppose that $g: G \rightarrow G$ is an irreducible train track map, that E_1 and E_2 are edges in G with a common terminal vertex and that $E_1 \bar{E}_2$ is an indivisible periodic Nielsen path that either has period one or is mapped by $f_\#$ to its inverse (and so has period two). Let G' be the graph obtained from G by isometrically identifying E_1 with E_2 and let $p: G \rightarrow G'$ be the quotient map. Then there is an induced irreducible train track map $g': G' \rightarrow G'$ with the following properties:*

- (1) $g'p = pg$.
- (2) If $f_g : \tilde{G} \rightarrow T_+$ and $f_{g'} : \tilde{G}' \rightarrow T_+$ are as in Corollary 2.14 then p lifts to $\tilde{p} : \tilde{G} \rightarrow \tilde{G}'$ such that $f_g = f_{g'}\tilde{p}$.
- (3) If g is a train track representative of $\phi \in \text{Out}(F_r)$ via a marking $\rho : R_r \rightarrow G$, then g' is also a train track representative of ϕ via the marking $p \circ \rho : R_r \rightarrow G'$.
- (4) The function p maps the principal vertices of g onto the principal vertices of g' , preserving Nielsen classes. This map is one-to-one except that the endpoints of $E_1\bar{E}_2$ are identified to a single point.

In this lemma, the local decomposition of $W(\phi)$ induced by g is obtained from that induced by g' by pasting the local stable Whitehead graphs of g' at the initial endpoints of E_1, E_2 , identifying the direction of E_1 with the direction of E_2 .

We shall refer to the construction of this lemma as a *Nielsen path collapse* performed on $E_1\bar{E}_2$.

Proof. The initial endpoints of E_1 and E_2 are distinct because ϕ does not fix the conjugacy class of any basis element. There is a path τ such that $g(E_i) = E_{\sigma(i)}\tau$ where $i = 1, 2$ and $\sigma : \{1, 2\} \rightarrow \{1, 2\}$. Denote $p(E_1) = p(E_2)$ by E' and define $g'(E') = E'p(\tau)$. Every other edge of G' equals $p(E)$ for some edge E of G and we define $g'(p(E)) = pg(E)$ which is an immersed path because p restricts to an immersion on legal paths in G . For the same reason $(g')^k(p(E)) = pg^k(E)$ for all $k > 0$ which implies that g' is an irreducible train track map. Item (1) is true by construction. Let $\tilde{p} : \tilde{G} \rightarrow \tilde{G}'$ be the lift that fixes all edges that do not project to E_1 or E_2 . Then f_g and $f_{g'}\tilde{p}$ are equivariant, agree on vertices and restrict to isometries on all edges. It follows that $f_g = f_{g'}\tilde{p}$, proving (2). Item (3) is obvious.

To prove (4), let w_i be the initial point of E_i and let $w' = pw_1 = pw_2$. The map p is injective on vertices except that $pw_1 = pw_2$. A vertex v of G is g -periodic if and only if pv is g' -periodic; for each of the two cases $v \in \{w_1, w_2\}$ and $v \notin \{w_1, w_2\}$ this is obvious, but for slightly different reasons. We now assume that $v \in G$ is a principal vertex and prove, by consideration of cases, that $pv \in G'$ is principal. Case 1: if v has at least three Dg -periodic directions then their images under $D_v p$ give at least three Dg' -periodic directions at pv . Case 2: if v is the endpoint of a periodic Nielsen path γ other than $E_1\bar{E}_2$ or its inverse, then $p\gamma = p_\# \gamma$ is a periodic Nielsen path and has pv as an endpoint. Case 3: if $v = w_1$ or w_2 then Dg has at least two periodic directions at each of w_1 and w_2 , and so $pv = w'$ has at least three periodic directions, namely $D_{w_1}p(E_1) = D_{w_2}p(E_2)$, one more in $\text{image}(D_{w_1}p)$, and one more in $\text{image}(D_{w_2}p)$. It follows that if v is principal then so is pv . Conversely, every vertex of G' with at least three periodic directions occurs as in Case 1 or Case 3, and every periodic Nielsen path of g' occurs as in Case 2. This proves that v is principal for g if and only if pv is principal for g' , and moreover that Nielsen classes are preserved. The statement about local decompositions is obvious by construction. \diamond

A Nielsen path collapse is similar to the more familiar operation from [BH92] of doing a fold on a turn of a train track map to produce a new train track map. We note some of the properties of the latter; the proof is similar to Lemma 4.1 but easier.

Lemma 4.2. *Consider an irreducible train track map $g: G \rightarrow G$ and oriented edges $E_1, E_2 \subset G$ with a common terminal vertex so that $g(E_1) = g(E_2)$. Let $p: G \rightarrow G'$ be the Stallings fold, the quotient map that identifies E_1 to E_2 . Then there is an induced irreducible train track map $g': G' \rightarrow G'$ that satisfies items (1)–(3) of Lemma 4.1, and such that*

- (4') *The function p restricts to a bijection between the principal vertices of g and g' preserving Nielsen classes. The local decompositions of the ideal Whitehead graph induced by g and by g' are identical.*

◇

We say that g' is obtained from g by a *fold* of E_1 with E_2 .

4.2 Cutting local stable Whitehead graphs

A Nielsen path collapse is reversible, and one can ask how to recognize when a given irreducible train track map $g': G' \rightarrow G'$ is the product of a Nielsen path collapse applied to some $g: G \rightarrow G$ as in Lemma 4.1. Since the map $p: G \rightarrow G'$ restricts to an immersion on leaves of Λ_- it induces a map of local Whitehead graphs $Dp: W_w \rightarrow W_{p(w)}$ for each vertex w of G . Similarly $g: G \rightarrow G$ induces a map $Dg: W_w \rightarrow W_{g(w)}$. If $w_i \in G$ is the initial vertex of E_i and $w' = p(w_i) \in G'$ then $W_{w'}$ decomposes as the union of $Dp(W_{w_1})$ and $Dp(W_{w_2})$ and these two subgraphs of $W_{w'}$ intersect in the single vertex determined by $E' = p(E_i)$. This vertex is therefore a *cut vertex* of $W_{w'}$, meaning that it separates $W_{w'}$. Note also that if $g'(v') = w'$ for some vertex $v' \neq w'$ then $Dg'(W_{v'})$ is contained in either $Dp(W_{w_1})$ or $Dp(W_{w_2})$. Finally, w' is a fixed point of g' and the action of Dg' on $W_{w'}$ preserves the decomposition $W_{w'} = Dp(W_{w_1}) \cup Dp(W_{w_2})$, fixing or interchanging the two sets $Dp(W_{w_1})$, $Dp(W_{w_2})$ depending on whether $g_\#$ preserves or reverses $E_1 \bar{E}_2$.

The following lemma states that these necessary conditions are also sufficient. The lemma is stated with sufficient breadth so that the construction may be used as well to reverse a fold.

Lemma 4.3. *Suppose that $g': G' \rightarrow G'$ is an irreducible train track map, that w' is a vertex of G' and that*

- (1) *$W_{w'}$ is the union of two non-trivial subgraphs X_1 and X_2 that intersect in a single vertex x .*
- (2) *If $g'(v') = w'$ for some vertex $v' \neq w'$ then $Dg'(W_{v'})$ is contained in either X_1 or X_2 .*
- (3) *If $g'(w') = w'$ then $Dg'(x) = x$ and there exists $\sigma: \{1, 2\} \rightarrow \{1, 2\}$ such that $Dg'(X_i) = X_{\sigma(i)}$.*

Then there exists an irreducible train track map $g: G \rightarrow G$ and a pair of edges E_1 and E_2 in G with a common terminal vertex and with distinct principal initial vertices w_1 and w_2 such that the quotient map $p: G \rightarrow G'$ that identifies E_1 to E_2 is a homotopy equivalence satisfying:

- (4) *$pg = g'p$.*
- (5) *$p(w_1) = p(w_2) = w'$.*

(6) $p(E_1) = p(E_2) = E'$ is the edge in $W_{w'}$ corresponding to x .

(7) $Dp(W_{w_i}) = X_i$ for $i = 1, 2$.

Furthermore, if $g'(w') = w'$ then g and the edges E_1, E_2 satisfy the hypotheses of Lemma 4.1 and g' is obtained from g by a Nielsen path collapse along $E_1 \bar{E}_2$, whereas if $g'(w') \neq w'$ then g and E_1, E_2 satisfy the hypotheses of Lemma 4.2 and g' is obtained from g by a fold of E_1 with E_2 .

Proof. For $i = 1, 2$ denote the edges of G' that correspond to vertices of X_i other than x by P'_i . Let u' be the terminal vertex of the edge E' corresponding to x . Define a new graph G from G' by replacing w' with a pair of vertices w_i , by replacing E' with a pair of edges E_i with initial vertex w_i and terminal vertex u' and by reattaching the edges in P'_i to w_i .

The map $p : G \rightarrow G'$ that identifies E_1 to E_2 is a homotopy equivalence. We construct $g : G \rightarrow G$ as follows. If $g'(w') = w'$ then $g(w_i) := w_{\sigma(i)}$; otherwise $g(w_i) := p^{-1}g'(w')$. If v is a vertex of G such that $g'(p(v)) \neq w'$, then $g(v) := p^{-1}(g'(p(v)))$. If v is neither w_1 nor w_2 and if $g'(p(v)) = w'$ then $g(v) := w_i$ where $Dg'(W_v) \subset X_i$. If E is an edge with endpoints s and t then define $g(E)$ to be the unique path with endpoints $g(s)$ and $g(t)$ such that $pg(E) = g'p(E)$. This unique lifting exists because the initial and terminal vertices have already been lifted and the base vertex of each turn in $g'(p(E))$ lifts uniquely: at the only vertex w' that does not lift uniquely to G the turns of $g'(p(E))$ lift uniquely because they are contained in either X_1 or X_2 . This completes the definition of g . Properties (4)–(7) are clear from the construction.

We must still show that $g : G \rightarrow G$ is an irreducible train track map. The obvious induction argument shows that $pg^k = (g')^k p$ for all $k \geq 0$. If E is an edge of G then $p(E)$ is an edge of G' and so $(g')^k p(E)$ is an immersed path for all $k \geq 0$. It follows that $g^k(E)$ is immersed and hence that g is a train track map. Each $(g')^k(pE)$ is a subpath of a leaf of the realization of Λ_- in G' . It follows that $g^k(E)$ is a subpath of a leaf of the realization of Λ_- in G . Example 2.5(1) of [BFH97] implies that $g^k(E)$ covers G for all sufficiently large k and hence that g is irreducible.

It remains to prove the “furthermore” clause. If $g'(w') = w'$ then clearly $E_1 \bar{E}_2$ is a periodic Nielsen path mapped by $g_\#$ to itself or its inverse, so the Nielsen collapse of $E_1 \bar{E}_2$ is defined and its product is clearly g' . If $g'(w') \neq w'$ then $g(w_1) = g(w_2) = p^{-1}(g'(w'))$, and clearly $g(E_1) = g(E_2)$ is the unique path in G with initial point $p^{-1}(g'(w'))$ that projects to $g'(E)$; the fold of E_1 with E_2 is therefore defined and its product is clearly g' . \diamond

4.3 The finest local decomposition

Here is our main result, in which we describe a process that starts with an arbitrary train track representative and produces one whose local decomposition is finer than any other. We do this only in the rotationless case; a more general result, while interesting, is not necessary for our present purposes, and the technical details would be a distraction.

Proposition 4.4. *Let $\phi \in \text{Out}(F_r)$ be nongeometric, fully irreducible, and rotationless. If $g : G \rightarrow G$ is an irreducible rotationless train track map representing ϕ with the maximal number of principal vertices then SW_v has no cut vertices for any principal vertex v of g , and the local decomposition of $W(\phi)$ induced by g is the cut point decomposition.*

The proof gives a process that will produce such a g starting from any train track representative of ϕ . The steps in the process involve some folds, some inverse Nielsen path collapses, and some inverse folds.

Proof. We begin with some preliminary folding that does not change the number of principal vertices. Recall that a pair of directions d_1 and d_2 at the same vertex w are in the same gate at w if they are identified by some iterate of Dg . Define the *gate excess* at w to be the number of directions minus the number of gates. Define the *total gate excess* to be the sum of the gate excesses over all principal vertices and over all nonprincipal, and hence non-periodic, vertices with three or more gates. If w has two directions d_1, d_2 that are identified by Dg itself then, after possibly subdividing G , we can fold the edges corresponding to d_1 and d_2 to create a new nonprincipal vertex with two gates, and so the total gate excess is reduced as long as w is not a nonprincipal vertex with two gates. Since g is rotationless, this can be applied repeatedly to the principal vertices until each of them has zero gate excess. If the total gate excess is still positive then there exists a nonprincipal vertex w with at least three gates and positive gate excess. It follows that some positive iterate of g takes w to a principal vertex, and so there exists a minimal $i \geq 1$ such that $g^i(w)$ has zero gate excess. Replacing w by $g^{i-1}(w)$ we may assume that there are two directions d_1, d_2 at w that are identified by Dg itself, and so a subdivision and fold at w reduces the total gate excess. Continuing by induction, we may assume that g has zero gate excess. In particular, for every vertex w which is either principal or has at least three gates, $Dg \mid W_w : W_w \rightarrow W_{g(w)}$ is injective. Furthermore, $W_w = SW_w$ if w is principal.

Assuming now that g has a principal vertex w such that $W_w = SW_w$ has a cut vertex, we describe a sequence of inverse folds followed by an inverse Nielsen collapse which refines the local decomposition of $W(\phi)$ induced by g . Let X_1 and X_2 be non-trivial subgraphs of W_w whose union is W_w and whose intersection is $\{x\}$. Define V to be the set of vertices $v \neq w$ for which there exists $i \geq 0$ so that $g^i(v) = w$ and so that $Dg^i(W_v)$ is not contained in either X_1 or X_2 . Thus $Dg^i(W_v)$ contains x , at least one vertex in $X_1 \setminus \{x\}$ and at least one vertex in $X_2 \setminus \{x\}$. In particular, v has at least three gates. Let $X'_1 \cup X'_2$ be the decomposition of W_v defined by pulling back the decomposition $X_1 \cup X_2$ of W_w via Dg^i . It follows, by injectivity of $Dg \mid W_v$, that $X'_1 \cap X'_2$ is a cut vertex x' . If u is a vertex such that $g(u) = v$ and such that $Dg(W_u)$ is not contained in either X'_1 or X'_2 then $u \in V$.

Case 1: $V = \emptyset$. In this case the three items in the hypothesis of Lemma 4.3 are satisfied: (1) is obvious, (2) follows from the definition of V and (3) follows from the assumption that ϕ is rotationless. Applying the lemma, an inverse Nielsen collapse produces an irreducible train track map with more principal vertices than g , refining the local decomposition of $W(\phi)$ induced by g .

Case 2: $V \neq \emptyset$. Since g is rotationless, there exists $v \in V$ that is not $g(u)$ for any $u \in V$. Then Lemma 4.3 applies to the decomposition $X'_1 \cup X'_2$ described above: (1) is obvious, (2) follows from the choice of v and our discussion above, and (3) is vacuous. Applying the lemma, an inverse fold produces a new irreducible train track map with the same number of principal vertices and the same local decomposition as $g : G \rightarrow G$ but with one fewer element in V . Continuing inductively until $V = \emptyset$, we reduce to Case 1. \diamond

5 Weak train tracks

Let $\phi \in \text{Out}(F_r)$ be nongeometric and fully irreducible with expanding lamination $\Lambda_- = \Lambda_-(\phi)$ and attracting F_r tree $T_+ = T_+(\phi)$. Consider a marked graph G or its universal cover $\tilde{G} = T$. We say that G or T is a *train track* for ϕ if there exists an irreducible rotationless train track map $g: G \rightarrow G$ representing an iterate (a positive power) of ϕ , and we say that G or T *arises from* $g: G \rightarrow G$. We say that G or T is a *weak train track* for ϕ if, up to renormalizing T and/or T_+ , there exists a Λ_- isometry $f_T: T \rightarrow T_+$. The definitions of train tracks and weak train tracks are evidently invariant under the equivalence relation defining \mathcal{X}_r , and so we can speak about a point of \mathcal{X}_r being a train track or a weak train track.

Since the definition of a weak train track for ϕ depends only on the lamination Λ_- and the tree T_+ , and since these two objects are insensitive to iterating ϕ , a weak train track for ϕ is the same thing as a weak train track for any iterate of ϕ . We will therefore replace ϕ with an iterate whenever it is convenient.

Henceforth we shall always fix a normalization of T_+ a priori. Having done so, we would like to normalize each weak train track T so that there exists a Λ_- isometry $f_T: T \rightarrow T_+$. As we will see in Section 5.4, there are examples where f_T is *not* uniquely determined by T , although the failure of uniqueness is mild (if Λ_- is nonorientable then f_T is unique; whereas if Λ_- is orientable then f_T varies over a space of maps homeomorphic to a closed interval in the real line). It might seem that the normalization of T depends on the choice of f_T . The following says that this does not happen:

Lemma 5.1. *Fix a normalization of T_+ . For each weak train track $T \in \mathcal{X}_r$ there exists a unique normalization of T in $\hat{\mathcal{X}}_r$ for which there exists a Λ_- isometry from T to T_+ .*

Proof. Consider any normalization of T for which there exists a Λ_- isometry $f_T: T \rightarrow T_+$. Consider the marked graph $G = T/F_r$. Each leaf L of Λ_- is realized by a surjective local isometry $\ell: \mathbf{R} \rightarrow G$, and so there is an oriented edge e of G and points $r < s \in \mathbf{R}$ such that $\ell(r) = \ell(s) \in \text{int}(e)$ and ℓ is orientation preserving near r and near s . Surgery of ℓ produces a loop by restricting ℓ to $[r, s]$ and identifying r and s . If $\gamma \in F_r$ represents this loop then the length of this loop is equal to the translation length of γ acting on T . Moreover, by construction the axis of γ in T is legal with respect to the map f_T , and so is mapped by f_T isometrically and γ -equivariantly onto the axis of γ in T_+ . The translation lengths of γ acting on T and on T_+ are therefore equal, thereby uniquely determining the normalization of T . \diamond

By Lemma 2.20, every train track for ϕ is a weak train track for ϕ . To be more explicit, for every train track representative $g: G \rightarrow G$ of ϕ , the train track \tilde{G} is a weak train track, and the map $f_g: \tilde{G} \rightarrow T_+$ is a Λ_- isometry. The main technical result of this section is Proposition 5.4 which explains how to obtain every sufficiently short weak train track by folding one particular train track. We also explain in Proposition 5.5 how to recognize when a weak train track is a train track. As an application we prove in Corollary 5.6 that the set of weak train tracks for ϕ is dense in the set of train tracks for ϕ , as subsets of \mathcal{X}_r .

Much of what we do in this section is to generalize some of the structure of train tracks, such as the local stable Whitehead graphs and the local decomposition of the ideal Whitehead graph $W(\Phi)$, to the context of a weak train tracks. These notions will depend not only on the weak train track T but also the choice of a Λ_- isometry $f_T: T \rightarrow T_+$. We will sometimes

emphasize the dependence on f_T by referring, for example, to the local decomposition of $W(\Phi)$ induced by T relative to f_T . As we proceed we will begin to suppress this dependence when f_T is clear from the context. On first pass, the reader may wish to assume uniqueness of f_T (a perfectly safe assumption when Λ_- is nonorientable).

5.1 Local decomposition of the ideal Whitehead graph.

Let T be a weak train track for ϕ . Choose a Λ_- isometry $f_T: T \rightarrow T_+$. If b is a branch point of T_+ and $\Phi \in \text{PA}(\phi)$ corresponds to b as in Lemma 2.16, then the realization in T of each leaf $\tilde{\ell} \in L(\Phi)$ contains a unique point w that is mapped by f_T to b , called the *base point* of $\tilde{\ell}$ in T . Fixing $w \in T$, as $\tilde{\ell}$ varies over the leaves in $L(\Phi)$ with base point w , the point w divides the realization of $\tilde{\ell}$ into two rays called *singular rays of T at w* . Each singular ray of T at w is mapped by f_T to a singular ray of T_+ at b . There are only finitely many F_r -orbits of basepoints w so we may equivariantly subdivide T to make each basepoint a vertex. A vertex of T is *principal* if it is the basepoint of some singular leaf. Base points of singular leaves, singular rays, and principal vertices are all defined relative to f_T .

For each point v of T , the *local Whitehead graph* $W(v; T)$ is the graph whose vertices are directions at v and with an edge connecting two vertices if the turn that they determine is taken by a leaf of $\tilde{\Lambda}_-$. Local Whitehead graphs are defined *absolutely*, not relative to f_T .

For each principal vertex w of T relative to f_T , the *local stable Whitehead graph* $SW(w; T)$ is the graph whose vertices are initial directions of singular rays based at w and with an edge connecting two vertices if the turn that they determine is taken by a singular leaf of Λ_- realized in T with base point w . If $f_T(w) = b$ then f_T induces an injection of $SW(w; T)$ into $W(b; T_+)$; we use this and Lemma 3.4 to identify $SW(w; T)$ with a subgraph of $W(b; T_+)$ and with a subgraph of the component $W(\Phi)$ of the ideal Whitehead graph $W(\phi)$ where $\Phi \in \text{PA}(\phi)$ corresponds to b . The local stable Whitehead graph $SW(w; T)$ is defined relative to f_T , so a more detailed notation which we prefer to avoid would be $SW(w; f_T)$.

The following lemma generalizes Lemma 3.2.

Lemma 5.2. *Suppose that T is a weak train track, $f_T: T \rightarrow T_+$ is a Λ_- isometry, b is a branch point of T_+ , $\Phi \in \text{PA}(\phi)$ is the representative of ϕ corresponding to b , and $\{w_i\}$ are the principal vertices of T that are mapped by f_T to b . Then*

- (1) $W(\Phi) = \cup SW(w_i; T)$.
- (2) *For $i \neq j$, $SW(w_i; T) \cap SW(w_j; T)$ intersect in at most one vertex. If they do intersect at a vertex P , then P is a cut point of $W(\Phi)$, in fact P separates $SW(w_i; T)$ from $SW(w_j; T)$ in $W(\Phi)$.*

Proof. (1) is immediate from the definitions and from our identification of $SW(w_i; T)$ with a subgraph of $W(\Phi)$. The proof of (2) is the same as the proof of the third item of Corollary 3.2. \diamond

As in the case of Corollary 3.2, Lemma 5.2 shows that the choice of a weak train track T and a Λ_- isometry $f_T: T \rightarrow T_+$ induces a decomposition of the ideal Whitehead graph $W(\phi)$, expressing each component of $W(\phi)$ as a union of local stable Whitehead graphs of T pasted together along cut vertices. We refer to this as the *local decomposition* of $W(\phi)$ relative to f_T .

Given two weak train tracks S, T and Λ_- isometries $f_S: S \rightarrow T_+$, $f_T: T \rightarrow T_+$, we say that f_T *splits at least as much as* f_S if the local decomposition $W(\phi) = \cup SW(v_j; T)$ is at least as fine as the local decomposition $W(\phi) = \cup SW(w_i; S)$, meaning that for each principal vertex v_j of T there is a principal vertex w_i of S such that $SW(v_j; T) \subset SW(w_i; S)$ where the inclusion takes place in $W(\phi)$.

To put it another way, a Λ_- isometry $f_T: T \rightarrow T_+$ imposes an equivalence relation on the set of singular leaves realized in T , where $l_1 \stackrel{f_T}{\sim} l_2$ if l_1, l_2 have the same base points relative to f_T . Obviously $l_1 \stackrel{f_T}{\sim} l_2$ if and only if their edges belong to the same component of the local decomposition of $W(\phi)$ relative to f_T . Thus, to say that f_T splits at least as much as f_S means that, for all singular leaves l_1, l_2 , if the realizations of l_1, l_2 in T have the same basepoint relative to f_T then their realizations in S have the same basepoint relative to f_S .

Lemma 5.3. *There exists a rotationless irreducible train track map $g: G \rightarrow G$ representing an iterate of ϕ such that for every weak train track T and Λ_- isometry $f_T: T \rightarrow T_+$, the Λ_- isometry $f_g: \tilde{G} \rightarrow T_+$ splits at least as much as f_T . If ϕ is rotationless then we may choose $g: G \rightarrow G$ to represent ϕ .*

Proof. After replacing ϕ with an iterate we may assume that ϕ is rotationless. Apply Lemma 4.4 to choose a rotationless train track map $g: G \rightarrow G$ whose local stable Whitehead links have no cut vertices. For each $\Phi \in \text{PA}(\phi)$, there are decompositions $W(\Phi) = \cup SW(\tilde{v}_i; \tilde{G})$ and $W(\Phi) = \cup SW(w_j; T)$ where $f_{\tilde{G}}(\tilde{v}_i) = f_T(w_j) = b$. Since $SW(\tilde{v}_i; \tilde{G})$ has no cut vertices, it must be contained in some single $SW(w_j; T)$. \diamond

5.2 Folding up to a weak train track

We now come to the main technical result of this section. This result says that from any fixed train track we can fold to any sufficiently short weak train track which splits no more than the fixed train track. This is very similar to Lemma 3.16 of [LL04] which combines results from [Lus99].

Proposition 5.4. *Let $\phi \in \text{Out}(F_n)$ be nongeometric and fully irreducible. For any train track representative $g: G \rightarrow G$ for ϕ with associated Λ_- isometry $f_g: \tilde{G} \rightarrow T_+$, there exists $\epsilon > 0$ so that if T is any weak train track for ϕ , if $f_T: T \rightarrow T_+$ is any Λ_- isometry, if $f_g: \tilde{G} \rightarrow T_+$ splits at least as much as f_T , and if $\text{Len}(T) \leq \epsilon$, then there exists a unique equivariant edge isometry $h: \tilde{G} \rightarrow T$ such that $f_g = f_T h$. This map h is moreover a Λ_- isometry.*

In this statement, recall our convention of fixing a normalization of T_+ , relative to which \tilde{G} and T are normalized, and so $\text{Len}(T)$ is well defined.

Proof. Uniqueness of h , and the fact that h must be a Λ_- isometry, are proved as follows. For any leaf L of Λ_- realized in \tilde{G} , since $f_g \mid L$ is an isometry, $h \mid L$ must be an injection, and since h is an edge isometry the map $h \mid L$ must be an isometry onto its image. Moreover, the image $h(L)$ must be the realization of L in T , which we denote L_T , proving that h is a Λ_- isometry. Since $f_T \mid L_T$ is an isometry, it follows that $h \mid L = (f_T \mid L_T)^{-1} \circ (f_g \mid L)$. This being true for any L , the map h is uniquely characterized.

For the rest of the proof we turn to existence of h .

There is no loss in assuming that ϕ is rotationless. After replacing $g: G \rightarrow G$ by an iterate we may assume that for any vertex v and any direction d at $v \in G$, the vertex $g(v)$ is fixed and the direction $D_v g(d)$ is fixed. We now assume the notation of Figure 2, and we write \tilde{G}_0 for \tilde{G} and f_{g_0} for f_g .

The map $h: \tilde{G}_0 \rightarrow T$ will be defined equivariantly, first on the vertices of \tilde{G}_0 , and then by linear extension over the edges. With appropriate choice of ϵ we will then prove that $f_{g_0} = f_T h$, from which it will immediately follow that h is edge isometric.

If $\tilde{x} \in \tilde{G}_0$ projects to a fixed point of g then there is a natural choice of $h(\tilde{x}) \in T$. To take care of a general vertex $\tilde{x} \in \tilde{G}_0$, we first push \tilde{x} forward via $\tilde{g}_{1,0}$ from \tilde{G}_0 to a point $\tilde{y} \in \tilde{G}_1$ which does project to a fixed point, and which now has a natural choice of image in T . In other words, for each vertex \tilde{x} of \tilde{G}_0 , we shall define $h(\tilde{x}) = h_1 \tilde{g}_{1,0}(\tilde{x})$ where h_1 is defined on $\tilde{y} = \tilde{g}_{1,0}(\tilde{x}) \in \tilde{G}_1$ as follows.

If \tilde{y} is principal then it is the basepoint of the realization of some singular leaf $\tilde{\ell}$ in \tilde{G}_1 and we define $h_1(\tilde{y})$ to be the basepoint of the realization of $\tilde{\ell}$ in T . This is independent of the choice of $\tilde{\ell}$, and therefore well defined, by the hypothesis that \tilde{G}_0 , and so also \tilde{G}_1 , splits at least as much as T . If \tilde{y} is not principal then there are exactly two fixed directions at \tilde{y} and they determine a leaf $\tilde{\ell}$ of $\tilde{\Lambda}_-$ by Lemma 2.17. There is a lift of g that fixes \tilde{y} and if Φ is the automorphism corresponding to that lift then $\Phi_{\#}(\tilde{\ell}) = \tilde{\ell}$. The homeomorphism $\Phi_+ : T_+ \rightarrow T_+$ of Corollary 2.21 preserves $\tilde{\ell}$, expanding the metric by $\lambda(\phi)$, and so has a unique fixed point c in the realization of $\tilde{\ell}$ in T_+ . Define $h_1(\tilde{y})$ to be the unique point in the realization of $\tilde{\ell}$ in T that is mapped to c by f_T .

Having defined h on the vertices of \tilde{G}_0 , we now extend it linearly to all of \tilde{G}_0 . Choose $\epsilon > 0$ to be less than half the length of the shortest edge in G_1 . We show next that for any edge \tilde{E} of \tilde{G}_0 , the path $h(\tilde{E})$ is legal with respect to f_T .

Let \tilde{x}_0 and \tilde{x}_1 be the endpoints of \tilde{E} . For $i = 0, 1$, let $\tilde{y}_i = \tilde{g}_{1,0}(\tilde{x}_i)$. There is a lift of g that fixes \tilde{y}_i and that fixes the direction at \tilde{y}_i determined by $\tilde{g}_{1,0}(\tilde{E})$; let Φ_i be the automorphism corresponding to this lift of g . We may choose the leaf $\tilde{\ell}_i$ used in the definition of $h_1(\tilde{y}_i)$ to ‘contain’ the direction at \tilde{y}_i determined by $\tilde{g}_{1,0}(\tilde{E})$. To be precise, $\tilde{\ell}_i$ is the leaf through \tilde{y}_i whose realization in \tilde{G}_1 is fixed by $(\Phi_i)_{\#}$. The leaf $\tilde{\ell}_i$ contains a non-trivial segment of $\tilde{g}_{1,0}(\tilde{E})$ incident to \tilde{y}_i , and its realization in T contains $h_1(\tilde{y}_i) = h(\tilde{x}_i)$. Let $\tilde{\ell}_2$ be any leaf whose realization in \tilde{G}_0 contains \tilde{E} , and so the realization of $\tilde{\ell}_2$ in \tilde{G}_1 contains $\tilde{g}_{1,0}(\tilde{E})$.

Let \tilde{z}_0, \tilde{z}_1 be the midpoints of the initial and terminal edges of $\tilde{g}_{1,0}(\tilde{E})$ respectively. For $i = 0, 1$, the realizations of $\tilde{\ell}_i$ and $\tilde{\ell}_2$ in T_+ contain an interval in T_+ of radius greater than ϵ centered on $f_{g_1}(\tilde{z}_i)$. Since T has length at most ϵ , the number ϵ is a cancellation constant (Lemma 2.5) for $f_T: T \rightarrow T_+$. It follows that the intersection of the realizations of $\tilde{\ell}_i$ and $\tilde{\ell}_2$ in T contains an open arc about a point \tilde{w}_i that maps by f_T to $f_{g_1}(\tilde{z}_i)$. One can therefore travel along a path in T that starts from $h(\tilde{x}_0) = h_1(\tilde{y}_0)$, goes along the realization of $\tilde{\ell}_0$ to \tilde{w}_0 , then along the realization of $\tilde{\ell}_2$ to \tilde{w}_1 , and then along the realization of $\tilde{\ell}_1$ to $h_1(\tilde{y}_1) = h(\tilde{x}_1)$, without ever taking an illegal turn. This legal path is therefore the unique embedded arc in T from $h(\tilde{x}_0)$ to $h(\tilde{x}_1)$, which by definition is $h(\tilde{E})$.

By construction, $f_{g_0}(x) = f_T h(\tilde{x})$ for all vertices \tilde{x} of \tilde{G}_0 . Thus $f_T h: \tilde{G}_0 \rightarrow T_+$ is an equivariant map that agrees with f_g on vertices and is linear on edges. It follows that $f_{g_0} = f_T h$ as desired. Also, since f_{g_0} and f_T are both isometric on edges, it follows that h

is isometric on edges.

◇

5.3 Comparing train tracks to weak train tracks

There are two moves one can perform on a train track map $g: G \rightarrow G$ representing an iterate of ϕ to produce a weak train track that is not a train track. The first move is to partially fold a pair of edges to create a vertex that is not preperiodic for the induced map. In this case, although there is an induced homotopy equivalence on the quotient space, this map does not take vertices to vertices no matter how the quotient space is subdivided as a CW complex. The second move is to collapse a pre-Nielsen path by a folding operation that is analogous to the one used in Lemma 4.1 without also collapsing the Nielsen path to which it maps. In this case there is no induced map on the resulting quotient space. The next proposition shows that these are the only possibilities, and the following corollary uses this to prove that train tracks are dense in the set of weak train tracks.

We say that $c \in T_+$ is a *periodic point* if there is an automorphism Φ representing a rotationless iterate of ϕ such that $c \in \text{Fix}(\Phi_+)$. Every branch point of T_+ is a periodic point by Lemma 2.16. Also, if $g: G \rightarrow G$ is a train track representative of some rotationless iterate of ϕ then Theorem 2.15 guarantees that if $x \in \tilde{G}$ projects to a periodic point of g then $f_g(x)$ is a periodic point of T_+ , proving the necessity of item (1) in the following:

Proposition 5.5. *Let $T = \tilde{G}$ be a weak train track and $f_T: T \rightarrow T_+$ a Λ_- isometry. The following are necessary and sufficient conditions for the existence of a train track representative $g: G \rightarrow G$ of ϕ^i for some i , such that $f_T = f_g$:*

- (1) *For every vertex w of T , $f_T(w) \in T_+$ is a periodic point.*
- (2) *For every vertex y of T , if $f_T(y) = b$ is a branch point of T_+ then there exists a principal vertex w of T such that $Df_T(W(y; T)) \subset Df_T(SW(w; T)) \subset W(b; T_+)$.*

Proof. For the necessity of item (2), choose $i \geq 1$ so that g^i maps every vertex to a fixed point of g^i and Dg^i maps every direction at a vertex to a fixed direction. Referring to the notation of Figure 2, it follows that there is a principal vertex w of $T = \tilde{G} = \tilde{G}_0$ such that y and w have the same image $w_i \in \tilde{G}_i$ under $\tilde{g}_{i,0} = \tilde{g}_{i,i-1} \circ \cdots \circ \tilde{g}_{1,0}$, and $Dg_{i,0}(W(y; T)) \subset SW(w_i; \tilde{G}_i)$. Since Dg^i is the identity on the set of fixed directions at fixed vertices of g^i , the map $D\tilde{g}_{i,0}$ induces an isomorphism between $SW(w; T)$ and $SW(w_i; \tilde{G}_i)$. Applying the “chain rule” $Df_T = Df_{g_0} = Df_{g_i} \circ Df_{g_{i,0}}$, it follows that $Df_T(W(y; T)) \subset Df_T(SW(w; T)) \subset W(b; T_+)$.

For sufficiency, let \mathcal{R} be the set of rotationless train track maps $g: G \rightarrow G$ representing an iterate of ϕ for which there exists an equivariant map $h: \tilde{G} \rightarrow T$ such that h is simplicial (vertices map to vertices and edges map to edges) and $f_g = f_T h$.

We first show that $\mathcal{R} \neq \emptyset$. By Lemma 5.3 there exists a rotationless train track map $g: G \rightarrow G$ representing an iterate of ϕ such that the associated train track $f_g: \tilde{G} \rightarrow T_+$ splits as least as much as $f_T: T \rightarrow T_+$. As there are only finitely many indivisible pre-Nielsen paths for g , we may assume, after replacing g with an iterate, that if μ is a pre-Nielsen path for g , then $g_\#(\mu)$ is a Nielsen path for g . Choose ϵ as in Proposition 5.4 applied to $f_g: \tilde{G} \rightarrow T_+$. Replacing G with $\phi^{-n}G$ has the effect of replacing the constant ϵ with $\lambda^n \epsilon$. Choosing n sufficiently large, we may therefore assume that $\text{Len}(T) \leq \epsilon$ and hence that there is an equivariant map $h: \tilde{G} \rightarrow T$ such that $f_{\tilde{G}} = f_T h$. Up to the action of F_r , there are only

finitely many points of \tilde{G} that map via h to vertices of T . Suppose that w is a vertex of T , that \tilde{x} is a point of \tilde{G} and that $h(\tilde{x}) = w$. Let x be the image of \tilde{x} in G . After replacing ϕ with an iterate, (1) and Theorem 2.15(1) imply the existence of a lift $\tilde{g}(x)$ of $g(x)$ such that $f_g(\tilde{g}(x)) = f_g(\tilde{x})$. Corollary 2.14 implies that for some $m \geq 0$ either $g^m(x)$ is the endpoint of Nielsen path or $g^{m+1}(x) = g^m(x)$. In either case $g^m(x) \in \text{Fix}(g)$. We may therefore subdivide along the g -orbit of x to make x and hence \tilde{x} a vertex. After finitely many such subdivisions, the map h is simplicial. We have now shown that $\mathcal{R} \neq \emptyset$.

For each $g: G \rightarrow G$ in \mathcal{R} the associated map $h: \tilde{G} \rightarrow T$ descends to a simplicial homotopy equivalence $\bar{h}: G \rightarrow H$ where the marked graph H is the quotient space of T by the action of F_r , and \bar{h} preserves marking. If \bar{h} is a homeomorphism then $T = \tilde{G}$ and $f_T = f_g$, and so we are done. Otherwise, there exist edges E_1 and E_2 of G , oriented so as to have a common terminal endpoint, such that E_1, E_2 are identified by the map \bar{h} . Arguing by cases that depend on the nature of E_1, E_2 , we will show that there is an element $g': G' \rightarrow G'$ of \mathcal{R} such that G' has fewer edges than G . The proof then concludes by induction.

There are two easy cases. In case 1, $E_1 \bar{E}_2$ is an indivisible Nielsen path and Lemma 4.1 produces the desired $g': G' \rightarrow G'$. In case 2, there exists $k \geq 1$ such that $g_{\#}^k(E_1) = g_{\#}^k(E_2)$, and replacing g by g^k , Lemma 4.2 produces the desired $g': G' \rightarrow G'$.

If neither case 1 nor case 2 applies to the pair E_1, E_2 , then Corollary 2.14 implies, after replacing g with an iterate if necessary, that $\sigma := g_{\#}(E_1 \bar{E}_2)$ is an indivisible Nielsen path. Lift σ to a path $\tilde{\sigma} \subset \tilde{G}$ with endpoints \tilde{v}_1 and \tilde{v}_2 . There are lifts \tilde{E}_i with a common terminal endpoints and with initial endpoints \tilde{x}_i such that (in the notation of Figure 2) $\check{g}_0(\tilde{x}_i) = \tilde{v}_i$. Let $y = h(\tilde{x}_1) = h(\tilde{x}_2)$ and $w_i = h(\tilde{v}_i)$. In particular, the endpoints v_1 and v_2 of σ are principal. Lemma 2.11 and Lemma 2.16 imply that $f_T(y) = f_g(\tilde{x}_i) = f_g(\tilde{v}_i) = f_T(w_i) = b \in T_+$ is a branch point. By construction, $Df_g(W(\tilde{x}_i; \tilde{G})) \subset Df_T(W(y; T))$ and $Df_g(W(\tilde{x}_i; \tilde{G})) \subset Df_g(SW(\tilde{v}_i; \tilde{G})) \subset Df_T(SW(w_i; T))$. For each $i = 1, 2$, the graph $Df_T(W(y; T))$ therefore intersects $Df_T(SW(w_i; T))$ in at least one edge. Applying (2) and Lemma 5.2 it follows that $Df_T(W(y; T)) \subset Df_T(SW(w_i; T))$, and so $Df_T(SW(w_1; T))$ and $Df_T(SW(w_2; T))$ intersect each other in at least one edge. Another application of Lemma 5.2 implies that $w_1 = w_2$. In particular, letting (E'_1, E'_2) be the illegal turn in σ , we have that $\bar{h}(E'_1) = \bar{h}(E'_2)$. Replacing E_1 and E_2 with E'_1 and E'_2 we are in case 1 if $\sigma = E'_1 \bar{E}'_2$ and in case 2 otherwise. \diamond

Corollary 5.6. *For any nongeometric, fully irreducible ϕ , a dense set of weak train tracks are train tracks.*

Proof. A point in T is said to be *preperiodic* if its image under f_T is a periodic point in T_+ .

We define a complexity $C(T)$ for weak train tracks T as follows. First define the ‘valence complexity’ $VC(T)$ to be the non-increasing sequence of integers with one entry k_i for each F_r -orbit of vertices w_i of valence $k_i \geq 3$. Then define the ‘aperiodic complexity’ $AC(T)$ to be the number of orbits of vertices in T that are not preperiodic. Finally, define $C(T) := (VC(T), AC(T))$ ordered lexicographically. Thus $C(T') < C(T)$ if $VC(T') < VC(T)$ or if $VC(T') = VC(T)$ and $AC(T') < AC(T)$.

It suffices to show that if T is not a train track then there is an arbitrarily close weak train track T' with $C(T') < C(T)$.

Let $g: G \rightarrow G$ be a rotationless irreducible train track map representing an iterate of ϕ . In the tree \tilde{G} there is a dense set of points that project to periodic points in G , and so for any leaf $\tilde{\ell}$ of $\tilde{\Lambda}_-$ realized in T_+ the set of periodic points is dense in $\tilde{\ell}$. The set of preperiodic points in T is therefore dense.

Suppose that w is a vertex of T that is not preperiodic. Then $f_T(w)$ is not a branch point, and so the local Whitehead graph $W(f_T(w); T_+)$ is a single edge. Consider the map $Df_T: W(w; T) \rightarrow W(f_T(w); T_+)$. If w has valence three then there are directions $d_1 \neq d_2$ at w which have the same image under Df_T . Folding these directions up to nearby points whose f_T -image is periodic produces a new valence three vertex whose f_T -image is preperiodic and reduces w to valence two; we then remove w from the list of vertices of T . This does not change $VC(T)$ and lowers $AC(T)$. If the valence of w is higher than three then perform a similar small fold to reduce $VC(T)$.

We may therefore assume that each vertex of T is preperiodic. By this same folding argument we may assume that if the image $Df_T(W(w; T))$ is not a single edge of $W(f_T(w); T_+)$ then the restricted map $Df_T: W(w; T) \rightarrow W(f_T(w); T_+)$ is an injection.

Proposition 5.5 implies that if T is not a train track, then there is a vertex y of T such that $f_T(y) = b$ is a branch point of T_+ but $Df_T(W(y; T)) \subset W(b; T_+)$ is not contained in $Df_T(SW(w; T))$ for any principal vertex w of T . In particular, $Df_T(W(y; T))$ contains at least three vertices of $W(b; T_+)$, so $Df_T(W(y; T))$ is not a single edge of $W(f_T(w); T_+)$, and so Df_T is injective on $W(y; T)$. Lemma 5.2 implies that $W(y; T)$ has a cut vertex d , and so $W(y; T)$ is the union of two subcomplexes X_1 and X_2 , each containing d properly, such that $X_1 \cap X_2 = \{d\}$. We now perform a small splitting very similar to the one in the proof of Lemma 4.3. Let E be the edge in the link of y that corresponds to d and let P_i be the edges in the link of y that correspond to the vertices of $X_i \setminus \{d\}$. After subdividing T at preperiodic points, we may assume that the length of E is arbitrarily small. Define a new tree T' from T by equivariantly replacing y with a pair of vertices y_i , by equivariantly replacing E with a pair of edges E_i with initial vertex y_i and terminal vertex equal to the terminal vertex of E and by equivariantly reattaching the edges in P'_i to y_i . The map $f_{T'}$ defined in the obvious way makes T' a weak train track of lesser complexity. \diamond

5.4 Rigidity and irrigidity of Λ_- isometries

In this section we study uniqueness of Λ_- isometries. Given a marked graph $G \in \hat{\mathcal{X}}_r$ with universal cover $T = \tilde{G}$ let $I(T, T_+)$, denote the space of Λ_- isometries $T \mapsto T_+$ equipped with the topology of pointwise convergence. Assuming that T is a weak train track, in other words $I(T, T_+)$ is nonempty, we wish to know whether this space has more than one point, and if it does then we wish to have some description of its structure. The answer, as it turns out, depends on orientability properties of Λ_- : if Λ_- is nonorientable then $I(T, T_+)$ is a single point; and if Λ_- is orientable then $I(T, T_+)$ is a compact interval, possibly degenerating to a point.

Recall that $D^2\partial F_r$ is the set of ordered pairs of elements of ∂F_r , that $\mathcal{G}F_r$ is the space of unordered pairs of elements of ∂F_r and that Λ_- is a closed F_r -equivariant subset of $\mathcal{G}F_r$. An *orientation* of Λ_- is a lift of the inclusion map $\Lambda_- \hookrightarrow \mathcal{G}F_r$ to a continuous, F_r -equivariant map $o: \Lambda_- \rightarrow D^2\partial F_r$. Equivalently, an orientation of Λ_- is a choice of orientation on each leaf of Λ_- with the following consistency condition: for some, and hence every, F_r tree T ,

there is a constant C so that if L_T and L'_T are the realizations in T of oriented leaves L and L' of Λ_- and if $L_T \cap L'_T$ has length at least C then the orientations of L_T and L'_T agree on $L_T \cap L'_T$. If Λ_- has an orientation then it is *orientable*. Since Λ_- is minimal, there are covering translations $\gamma \in F_r$ such that $L'_T \cap \gamma \cdot L_T$ is arbitrarily long. It follows that an orientation on Λ_- is determined by an orientation of any of its leaves and hence that Λ_- has either zero or two orientations.

An *orientation of Λ_- in T* is a choice of orientation on each leaf of Λ_- with the feature that the orientations on L_T and L'_T agree on $L_T \cap L'_T$ whenever $L_T \cap L'_T$ contains a nontrivial interval. An orientation of Λ_- in T therefore induces a well defined F_r -equivariant orientation on each edge of T . Recall from Lemma 2.20 that each leaf of Λ_- has a well defined realization in T_+ . Thus this definition works equally well to define an *orientation of Λ_- in T_+* .

It is obvious that if Λ_- is orientable in some T then it is orientable. We have the following partial converse.

Lemma 5.7. *If Λ_- is orientable then it is orientable in T_+ and is orientable in any weak train track T .*

Proof. Let $g : G \rightarrow G$ be a train track representative of ϕ , let \tilde{G} be the universal cover of G and let o be an orientation on Λ_- . Choose a lift $\check{g} : \tilde{G} \rightarrow \tilde{G}$ of g . Since $\hat{g} : \partial F_r \rightarrow \partial F_r$ is continuous and F_r -equivariant, it acts on the set of orientations of Λ_- and this action is independent of the choice of lift \check{g} . After replacing g with g^2 if necessary, we may assume that \hat{g} preserves o . Given L and L' so that $L_{\tilde{G}} \cap L'_{\tilde{G}}$ is a nontrivial interval, choose $n > 0$ so that $\check{g}^n(L'_{\tilde{G}}) \cap \check{g}^n(L_{\tilde{G}})$ has length greater than the constant C associated to o in the definition of an orientation. The orientations of $\check{g}^n(L'_{\tilde{G}})$ and $\check{g}^n(L_{\tilde{G}})$ agree on $\check{g}^n(L'_{\tilde{G}}) \cap \check{g}^n(L_{\tilde{G}})$ which implies that the orientations of $L'_{\tilde{G}}$ and $L_{\tilde{G}}$ agree on $L'_{\tilde{G}} \cap L_{\tilde{G}}$. This proves that Λ_- is orientable in \tilde{G} .

If $L_{T_+} \cap L'_{T_+}$ is a non-trivial interval then there is a train track representative $g : G \rightarrow G$ with universal cover \tilde{G} so that $L_{\tilde{G}}$ and $L'_{\tilde{G}}$ intersect in a non-trivial interval. By our previous argument, the orientations of $L_{\tilde{G}}$ and of $L'_{\tilde{G}}$ agree on $L_{\tilde{G}} \cap L'_{\tilde{G}}$. Since $f_g : \tilde{G} \rightarrow T_+$ preserves orientation on the leaves of Λ_- , it follows that the orientations of L_{T_+} and L'_{T_+} agree on $L_{T_+} \cap L'_{T_+}$. This proves that Λ_- is orientable in T_+ . If T is a weak train track then leaves of Λ_- that intersect in T also intersect in T_+ . An orientation of Λ_- in T_+ therefore pulls back to an orientation on Λ_- in T . \diamond

The next theorem gives an abstract description of the set $I(T, T_+)$. An orientation o of Λ_- induces orientations of Λ_- in T and in T_+ and each element $f \in I(T, T_+)$ preserves the induced orientations. It follows that for each vertex v of S , if e_1, e_2 are two edges at v in the same gate with respect to f then e_1, e_2 point in the same direction at v : either both point towards v in which case that gate is said to be a *negative* gate at v ; or both point away from v in which case the gate is said to be *positive*.

Theorem 5.8. *Fix a weak train track $G \in \hat{\mathcal{X}}_r$ with universal cover $T = \tilde{G}$, and fix $f \in I(T, T_+)$.*

- (1) *If Λ_- is not orientable in G then f is the unique element of $I(T, T_+)$.*
- (2) *If Λ_- is orientable in G then:*

- (a) *There exists a compact subinterval Δ of the real line, and a homeomorphism $\Delta \mapsto I(T, T_+)$. We denote this homeomorphism $t \mapsto f_t$.*
- (b) *For any leaf L of Λ_- and for any $x \in L_{T_+}$, the map $\Delta \mapsto T_+$ given by $t \mapsto f_t(x)$ is an orientation preserving isometry onto a segment of the leaf L_{T_+} .*
- (c) *Given $f \in I(T, T_+)$, the following hold:*
 - i. *f is an interior point of $I(T, T_+)$ if and only if every vertex of T has exactly one positive and one negative gate with respect to f .*
 - ii. *f is the positive (resp. negative) endpoint of $I(T, T_+)$ if and only if some vertex of T has more than one positive (resp. negative) gate with respect to f .*
 - iii. *f is the unique point of $I(T, T_+)$ if and only if there exist vertices $v, w \in T$ such that v has at least two positive gates with respect to f and w has at least two negative gates with respect to f .*

Even when Λ_- is orientable, this lemma leaves open the possibility that f is unique, that is, the interval Δ degenerates to a point. See below for examples where Δ is nondegenerate and where it is degenerate.

In item (2a) of this theorem, we use the topology of pointwise convergence on $I(T, T_+)$, but as should be clear from the proof this is equivalent to the Gromov topology on maps (for the definition of this topology see Section 6.4, proof of step 3(d)).

Proof. Consider two Λ_- isometric maps $f, f': T \rightarrow T_+$. For each oriented leaf L of Λ_- , let $L_T \subset T$ and $L_{T_+} \subset T_+$ denote the realizations of L as oriented, bi-infinite geodesics. Each of the restricted maps $f, f': L_T \rightarrow L_{T_+}$ is an orientation preserving isometry, and so these two maps differ by a real number $\delta(f, f'; L)$, meaning that the map $f': L_T \rightarrow L_{T_+}$ is obtained by postcomposing the map $f: L_T \rightarrow L_{T_+}$ by the translation of L_{T_+} that displaces each point by the amount $\delta(f, f'; L)$ (remember that L_{T_+} is oriented).

Next we study how $\delta(f, f'; L)$ depends on L . Consider another oriented leaf L' of Λ_- . If $L_T \cap L'_T = J$ is a nondegenerate segment then the restriction to J of the map $f: L_T \rightarrow L_{T_+}$ equals the restriction to J of the map $f: L'_T \rightarrow L'_{T_+}$, and the same is true for f' . It follows that if L_T and L'_T induce the same orientation on J then $\delta(f, f'; L) = \delta(f, f'; L')$, whereas if L_T and L'_T induce opposite orientations on J then $\delta(f, f'; L) = -\delta(f, f'; L')$. Note also that $\delta(f, f'; L)$ is invariant under the action of F_r on oriented leaves.

Suppose now that Λ_- is nonorientable in T . Then for any oriented leaf L of Λ_- there exists a covering transformation $g \in F_r$ such that $L_T \cap g \cdot L_T$ is a nondegenerate segment J and the orientations induced on J by L_T and $g \cdot L_T$ are opposite. It follows that

$$\delta(f, f'; L) = \delta(f, f'; g \cdot L) = -\delta(f, f'; L)$$

and so

$$\delta(f, f'; L) = 0$$

implying that $f|_{L_T} = f'|_{L_T}$. Since this is true for every leaf L , and since realizations of leaves of Λ_- cover T , it follows that $f = f'$.

Suppose now that Λ_- is orientable in T , and fix an orientation. If L, L' are any two oriented leaves then there exists $g \in F_r$ such that $g \cdot L_T \cap L'_T = J$ is a nondegenerate segment and the orientations on J induced by $g \cdot L_T$ and L'_T agree. It follows that

$$\delta(f, f'; L') = \delta(f, f'; g \cdot L) = \delta(f, f'; L)$$

This shows that $\delta(f, f'; L)$ is a constant independent of L , which we denote $\delta(f, f')$. The map f' is therefore determined by f and the constant $\delta(f, f')$: for each $x \in T$, pick any leaf L whose realization in T contains x , and then $f'(x)$ is obtained by displacing $f(x)$ along L_{T_+} by the amount $\delta(f, f')$.

Fix $f \in I(T, T_+)$. Let Δ be the set of all real numbers of the form $\delta(f, f')$ for some $f' \in I(T, T_+)$. Each $d \in \Delta$ equals $\delta(f, f')$ for a unique $f' \in I(T, T_+)$ and we define $f_d = f'$. To complete the proof of (2a) we must show that Δ is a closed, bounded interval; continuity follows by noticing that for any $d, d' \in \Delta$ and any $x \in T$, the distance in T between $f_d(x)$ and $f_{d'}(x)$ is exactly $|d - d'|$. Item (2b) follows from item (2a) and the definition of $\delta(f, f')$.

To show that Δ is connected, given a Λ_- isometry $f': T \rightarrow T_+$, and given a number d between 0 and $\delta(f, f')$ it suffices to exhibit a Λ_- -isometry $f'': T \rightarrow T_+$ such that $\delta(f, f'') = d$. For $x \in T$ the segment $\overline{f(x)f'(x)}$ in T_+ has length $\delta(f, f')$, and this segment is oriented with initial point $f(x)$ when $\delta(f, f') > 0$ and terminal point $f(x)$ when $\delta(f, f') < 0$. Now define $f''(x)$ to be the point on this segment at displacement d from $f(x)$.

Boundedness of Δ follows from the observation that there exist leaves $L, L' \subset \Lambda_-$ such that $L_T \cap L'_T \neq \emptyset$ and $L_{T_+} \cap L'_{T_+}$ is bounded.

The fact that Δ is closed follows by application of the Ascoli-Arzelà theorem together with the observation that the compact-open limit of a sequence of Λ_- isometries $T \mapsto T_+$ is a Λ_- isometry, and the fact that on $I(T, T_+)$ the compact open topology agrees with the pointwise convergence topology.

Item 2(c)iii follows from items 2(c)i and 2(c)ii, to whose proofs we now turn. Suppose first that every vertex $v \in T$ has a unique positive gate with respect to f , so there is a positive length oriented segment J_v in T_+ with initial endpoint $f(v)$ such that every edge in T with initial endpoint v has an initial subsegment mapping isometrically to J_v . There being only finitely many orbits of vertices, we can truncate to assume that all the segments J_v have the same length $\epsilon > 0$. We can then alter f to a new Λ_- isometry, by moving each vertex v along J_v to the positive endpoint of J_v , and extending by an isometry on each edge of T . Applying item (2b) it follows that f is not the positive endpoint of $I(T, T_+)$.

Conversely, suppose that there exists a vertex $v \in T$ having two or more positive gates, and so there exist edges e, e' of T each with initial endpoint v such that $f(e), f(e')$ have distinct germs at $f(v)$. Choose leaves $L, L' \subset \tilde{\Lambda}_-$ whose realizations in T contain e, e' , and hence each passes through v . The realizations of L, L' in T_+ , while passing through $f(v)$, are disjoint in the positive direction beyond $f(v)$. By item (2b) it follows that f is the positive endpoint of $I(T, T_+)$. Combining this argument with that of the last paragraph, and repeating these arguments with orientation reversed, items 2(c)i and 2(c)ii follow. \diamond

5.5 Examples of exceptional weak train tracks

We describe first a weak train track T for which the space of Λ_- isometries $I(T, T_+)$ is not just a single point. Then we describe a weak train track T that is not a train track.

Example: a weak train track T with nontrivial $I(T, T_+)$. Let a, b and c be generators of F_3 . Consider $\phi \in \text{Out}(F_3)$ represented by the following positive automorphism Φ , shown

together with its inverse:

$$\Phi: \begin{cases} a \mapsto bacaaca \\ b \mapsto baca \\ c \mapsto caaca \end{cases} \quad \Phi^{-1}: \begin{cases} a \mapsto \bar{b}\bar{a}\bar{b}\bar{a}\bar{c} \\ b \mapsto b\bar{a}b \\ c \mapsto \bar{c}\bar{a}b\bar{c}\bar{a}\bar{b}\bar{a}b \end{cases}$$

As realized on the three petalled rose, these formulas each define train track maps representing ϕ and ϕ^{-1} , respectively. The automorphism Φ being positive, its transition matrix equals the matrix of its abelianization. Direct calculation shows that no eigenvalue of the abelianization of ϕ is a root of unity, so ϕ is fully irreducible by Lemma 3.5. Direct calculation of the Perron-Frobenius eigenvalues of the transition matrices of Φ and Φ^{-1} also show that $\lambda(\phi) > 6$ and $\lambda(\phi^{-1}) < 5$, so ϕ is not geometric.

Let $g: G \rightarrow G$ be the affine train track representative of ϕ determined by the above formula for Φ . Let λ be the expansion factor, Λ_- the expanding lamination, and T_+ the attracting tree for ϕ . When the three petalled rose G is oriented so that its three oriented edges represent a, b and c then it induces an orientation on Λ_- in $T = \tilde{G}$. It is clear that at the unique vertex of G the map g has exactly two positive gates $\{a, b\}$ and $\{c\}$ and one negative gate $\{\bar{a}, \bar{b}, \bar{c}\}$. Since each of these gates contains a fixed direction, these are also the gates for $f_g: \tilde{T} \rightarrow T_+$. Theorem 5.8 item 2(c)ii implies that f_g is the positive endpoint of $I(T, T_+)$ but is not the negative endpoint.


Now we shall slide the map g backwards to locate the negative endpoint of $I(T, T_+)$. Looking at the formula for Φ we see that the g image of each edge ends with the path aca . Postcomposing Φ with the inner automorphism $w \mapsto (aca)w(aca)^{-1}$ we obtain

$$\Phi': \begin{cases} a \mapsto acabaca \\ b \mapsto acab \\ c \mapsto acaca \end{cases}$$

This formula defines a train track representative $g': G \rightarrow G$ of ϕ supported on the *same marked graph* G as the train track representative g . Notice now that g' has one positive gate $\{a, b, c\}$ and two negative gates $\{\bar{a}, \bar{c}\}$ and $\{\bar{b}\}$ at its unique vertex. As before each gate contains a fixed direction so these are also the gates of $f_{g'}$. Thus $f_{g'}$ is the negative endpoint of $I(T, T_+)$.

Example: a weak train track that is not a train track. We describe an example of a nongeometric, fully irreducible $\phi \in \text{Out}(F_7)$ and a weak train track H for ϕ that is not a train track for any iterate of ϕ . We shall also sketch the proof that H is not even the domain of a nonsimplicial train track representative of any iterate of ϕ .

Start with the graph G and train track map $g: G \rightarrow G$ given by



One easily shows by the method of Stallings folds that g is a homotopy equivalence. The only illegal turn of g being $\{E, D\}$, and the words \overline{ED} and \overline{DE} not occurring in the definition of g , it follows that g is a train track map (of course, any map of a directed graph which preserves direction is a train track map).

Let g represent $\phi \in \text{Out}(F_3)$ with respect to any marking of G . Lemma 3.5 applies to show that ϕ is fully irreducible, and by computing some train track \bar{g} for ϕ^{-1} one shows that the expansion factors of g and \bar{g} are not equal and hence ϕ is not geometric (in fact this outer automorphism is conjugate to the one in Section 3.4, as one can show by computing and collapsing a Nielsen path of g to get the train track map in Section 3.4).


Next we compute the local Whitehead graphs at p and q . The turns taken by the images of edges under g are

$$\{\overline{A}, F\}, \{\overline{D}, F\}, \{\overline{F}, A\}, \{\overline{F}, D\}, \{\overline{F}, E\}$$

The map Dg acts as follows: it takes $\{\overline{A}, F\}$ and $\{\overline{F}, E\}$ to $\{\overline{F}, D\}$ and fixes the latter; it takes $\{\overline{D}, F\}$ to $\{\overline{E}, D\}$ and fixes the latter; and it fixes $\{\overline{F}, A\}$. These are all the turns taken by iterates of edges under g , so the local Whitehead graphs are

$$\begin{array}{l} \text{At } p: \quad \overline{A} \text{ --- } F \text{ --- } \overline{D} \\ \\ \text{At } q: \quad \begin{array}{c} A \text{ --- } \overline{F} \text{ --- } D \text{ --- } \overline{E} \\ \quad \quad \quad \downarrow \\ \quad \quad \quad E \end{array} \end{array}$$

Note that D is a cut vertex of the local Whitehead graph at q . We construct H by a small splitting of G at q in the direction D , as was done in Lemma 4.3 and Corollary 5.6. The effect of this is to split the vertex q into two vertices q_1, q_2 . An initial segment of the edge D is split into two new edges D_1, D_2 , with the tail of D_i incident to q_i , and with the heads of D_1, D_2 incident to a new valence 3 vertex r . The remaining terminal segment of D becomes an edge D_3 from r to p . The other four ends $A, E, \overline{E}, \overline{F}$ of G at q are detached from q and reattached as follows: A, E , and \overline{F} are reattached to q_1 , and \overline{E} is reattached to q_2 .



To be precise about metric and marking, the edges lengths of A, E, F in H are the same as in G ; $\text{Len}_H(D_1) = \text{Len}_H(D_2) = \epsilon > 0$; and $\text{Len}_H(D_3) = \text{Len}_G(D) - \epsilon$. Let $h: H \rightarrow G$ be the edge isometry that is the “identity” in A, E, F , and that maps $D_i * D_3$ isometrically to D for $i = 1, 2$. The marking of G composed with a homotopy inverse of h defines the marking of H .

Consider the Λ_- isometry $f_g: \tilde{G} \rightarrow T_+$, and let $\tilde{h}: \tilde{H} \rightarrow \tilde{G}$ be an F_3 equivariant lift of h . To show that H is a weak train track we show that $f_h = f_g \circ \tilde{h}: \tilde{H} \rightarrow T_+$ is a Λ_- isometry, and to do this we must check that each leaf L of Λ_- realized in G lifts isometrically via \tilde{h} to the realization of L in \tilde{H} . This follows immediately from the observation that in the local Whitehead graph of G at q , the four edges \overline{AF} , \overline{EF} , \overline{DF} , \overline{DE} lifts to edges \overline{AF} , \overline{EF} , $D_1\overline{F}$ at q_1 and $D_2\overline{E}$ at q_2 .

We prove that f_h is the unique Λ_- isometry from $\tilde{H} \rightarrow T_+$, by applying Theorem 5.8. Since Λ_- is oriented we must check condition 2(c)iii in the theorem. Note that the directions \overline{A} and \overline{D} of G at p are in distinct negative gates, because \overline{A} and \overline{D} are mapped by Dg to \overline{E} and \overline{F} at q , and the latter are fixed by Dg ; it follows that any lift to \tilde{H} of the vertex p of H has two negative gates, containing lifts of the directions \overline{A} and \overline{D} respectively. Note also that in G the directions A and D at q are in distinct positive gates, because they are fixed by g ; it follows that any lift to \tilde{H} of q_1 has two positive gates, containing lifts of the directions A and D_1 respectively.

To prove that H is not a train track of any iterate of ϕ , by uniqueness of the Λ_- isometry f_h it suffices to prove that f_h does not satisfy item (2) of Proposition 5.5. Let $\tilde{p} \in \tilde{H}$ be a lift of $p \in H$, and so $\tilde{h}(\tilde{p})$ is a lift of $p \in G$. We know that there is a lift $\tilde{q} \in G$ of $q \in G$ such that $b = f_g(\tilde{h}(\tilde{p})) = f_g(\tilde{q})$ is a branch point of T_+ and Df_g maps the four stable directions of \tilde{G} at \tilde{q} , named $\{\overline{A}, \overline{F}, \overline{D}, \overline{E}\}$ to the four directions of T_+ at b , which shall have the same names. We also know that f_g maps $\tilde{h}(\tilde{p})$ taking the three directions $\{\overline{A}, \overline{F}, \overline{D}\}$ at $\tilde{h}(\tilde{p})$ to the three directions $\{\overline{F}, \overline{D}, \overline{E}\}$ at b . It follows that Df_h maps the three directions $\{\overline{A}, \overline{F}, \overline{D}\}$ at \tilde{p} to the three directions $\{\overline{F}, \overline{D}, \overline{E}\}$ at b . The only principal vertices in \tilde{H} that map to b via f_h are the lifts \tilde{q}_1, \tilde{q}_2 of q_1, q_2 such that $\tilde{h}(\tilde{q}_1) = \tilde{h}(\tilde{q}_2) = \tilde{q}$. However, the image under Df_h of the set of stable directions of \tilde{H} at \tilde{q}_i is the set $\{\overline{A}, \overline{F}, \overline{D}\}$ at b when $i = 1$, and the set $\{\overline{D}, \overline{E}\}$ when $i = 2$, and neither of these two sets contains $\{\overline{F}, \overline{D}, \overline{E}\}$, as is required by item (2) of Proposition 5.5.

We also sketch the proof that H is not the domain of any *nonsimplicial train track representative* of any iterate of ϕ , meaning a homotopy equivalence $g: H \rightarrow H$ that represents ϕ such that for each edge e of H and each $i \geq 1$ the map $g|_{\text{int}(e)}$ is a local homothety with stretch factor $\lambda(\phi)^i$. To prove this requires generalizing several results about train track representatives to the context of nonsimplicial train track representative. Suppose that there is a nonsimplicial train track representative $g: H \rightarrow H$ of some iterate of ϕ . We can lift g to define maps as denoted in Figure 2 (replacing the symbol G by H), and then Theorem 2.13 generalizes to produce T_+ as a direct limit, and to show that the unique Λ_- isometry $f_h: \tilde{H} \rightarrow T_+$ is the direct limit map denoted by f_g in Figure 2. Now the argument given in Proposition 5.5 for the necessity of item (2) generalizes in a straightforward manner. Finally, we have already proved that f_h fails to satisfy item (2).

6 Topology of the axis bundle

For this section fix a nongeometric, fully irreducible $\phi \in \text{Out}(F_r)$ with $\Lambda_- = \Lambda_-(\phi)$ and $T_+ = T_+(\phi)$, and fix a normalization of T_+ . Having so far proved the equivalence of two of the three proposed definitions of the axis bundle — the set of weak train tracks (2), and the closure of the set of train tracks (3) — in this section we work with definition (2), in terms of which we prove a version of Theorem 1.2.

Let $\mathcal{A} \subset \mathcal{X}_r$ be the set of all weak train tracks of ϕ . By Lemma 5.1, each $T \in \mathcal{A}$ has a normalization $\widehat{T} \in \widehat{\mathcal{X}}_r$ well-defined by requiring existence of a Λ_- isometry $\widehat{T} \mapsto T_+$, and so $\text{Len}(T) = \text{Len}(\widehat{T}/F_r) \in (0, +\infty)$ is well-defined. This map satisfies the equivariance condition

$$\text{Len}(T \cdot \phi^i) = \lambda(\phi)^{-i} \text{Len}(T) \quad \text{for any } i \in \mathbf{Z}$$

because from any Λ_- isometry $\widehat{T} \mapsto T_+$ we obtain a Λ_- isometry $\widehat{T} \cdot \phi^i \mapsto T_+ \cdot \phi^i = \lambda(\phi)^i T_+$ which, after renormalizing, becomes a Λ_- isometry $(\lambda(\phi)^{-i} \widehat{T}) \cdot \phi^i \mapsto T_+$.

Let $\widehat{\mathcal{A}} \subseteq \widehat{\mathcal{X}}_r$ be the set of normalized weak train tracks; projection $\widehat{\mathcal{X}}_r \mapsto \mathcal{X}_r$ restricts to a bijection $\widehat{\mathcal{A}} \mapsto \mathcal{A}$. Abusing notation we often use the same symbol to represent a weak train track and its normalization, although we will use the symbol $\widehat{}$ to maintain the distinction when needed.

Here is the main result of this section, a major step in the proof of Theorem 1.2:

Theorem 6.1. *The map $\text{Len}: \mathcal{A} \rightarrow (0, +\infty)$ is a proper homotopy equivalence.*

Outline. Section 6.1 proves basic topological facts about \mathcal{A} , including continuity of the normalization map $\mathcal{A} \mapsto \widehat{\mathcal{A}}$ and the resulting continuity of the length map $\text{Len}: \mathcal{A} \rightarrow (0, +\infty)$. Section 6.2 proves a technical result, strengthening the topological properties of $\widehat{\mathcal{A}}$, that will be useful in the proof of Theorem 6.1. Section 6.3 proves that $\text{Len}: \mathcal{A} \rightarrow (0, +\infty)$ is a proper map, using folding methods combined with Proposition 5.4. Section 6.4 proves Theorem 6.1 using Skora’s method [Sko] for studying homotopy types of spaces of trees (see also [Whi93], [Cla05], [GL05]).

6.1 Continuity properties of the normalized axis bundle

We prove here several facts about the set $\widehat{\mathcal{A}}$ and related facts about the length map.

Lemma 6.2.

- (1) \mathcal{A} is a closed subset of \mathcal{X}_r and $\widehat{\mathcal{A}}$ is a closed subset of $\widehat{\mathcal{X}}_r$.
- (2) The normalization map $\mathcal{A} \rightarrow \widehat{\mathcal{A}} \subset \widehat{\mathcal{X}}_r$ is a homeomorphism.
- (3) The map $\text{Len}: \mathcal{A} \rightarrow (0, +\infty)$ is continuous, surjective, and ϕ -equivariant, where ϕ acts on $(0, +\infty)$ by multiplying by $\frac{1}{\lambda}$.
- (4) The map $-\log \text{Len}: \mathcal{A} \rightarrow \mathbf{R}$ is also continuous, surjective and ϕ -equivariant, where ϕ acts on \mathbf{R} by adding $\log \lambda$.

Proof. Note that (4) follows from (3). Also, in proving (3), continuity of Len follows from (2) and equivariance follows from Theorem 2.15. For surjectivity, choose a train track representative $g: G \rightarrow G$, let $T_0 = \tilde{G}$, let $T_1 = T_0 \cdot \phi$, lift g to an F_r -equivariant Λ_- isometry $f: T_0 \rightarrow T_1$, and interpolate f by a Stallings fold path from T_0 to T_1 . For each T on this path there are equivariant edge isometries $T_0 \xrightarrow{f'} T \xrightarrow{f''} T_1$ that factor f , and so f'' followed by a Λ_- isometry $T_1 \rightarrow T_+$ is a Λ_- isometry $T \rightarrow T_+$, proving that the fold path is in $\hat{\mathcal{A}}$. Length varies continuously and monotonically along this path from T_0 to T_1 with image $[\text{Len}(G), \text{Len}(G)/\lambda]$. Surjectivity of $\text{Len}: \hat{\mathcal{A}} \rightarrow (0, +\infty)$ then follows from ϕ -equivariance.

Consider a sequence $G_i \in \mathcal{A}$ which converges in \mathcal{X}_r to some limit G . We first show that the normalized sequence \hat{G}_i converges to some limit $\hat{G} \in \mathcal{X}_r$, and second that $\hat{G} \in \hat{\mathcal{A}}$. This proves (1). It also proves that the normalization map $\mathcal{A} \rightarrow \hat{\mathcal{A}}$ is continuous, but its inverse is the restriction of the continuous projection map $\hat{\mathcal{X}}_r \mapsto \mathcal{X}_r$, proving (2) as well. Let T, T_i denote the universal covers of \hat{G}, \hat{G}_i .

Cells in \mathcal{X}_r being locally finite, by restricting to each of the finitely many closed cells containing G we may assume that all the G_i are in a single open cell, so there is a family of marked homeomorphisms $g_{ji}: G_i \rightarrow G_j$ such that $g_{kj} \circ g_{ji} = g_{ki}$, well-defined by requiring each g_{ji} to be affine on each edge of G_i . Lift g_{ji} to an F_r -equivariant homeomorphism $\tilde{g}_{ji}: T_i \rightarrow T_j$ which is affine on each edge.

Convergence in $\hat{\mathcal{A}}$ is defined by convergence of translation length functions ℓ_γ for all $\gamma \in F_r - \{\text{Id}\}$, whereas convergence in \mathcal{A} is defined by convergence of ratios $\ell_\gamma/\ell_{\gamma'}$ for all $\gamma, \gamma' \in F_r - \{\text{Id}\}$. Since G_i converges to G in \mathcal{A} , to prove convergence of \hat{G}_i it suffices to find a single $\gamma \in F_r - \{\text{Id}\}$ such that $\ell_\gamma(T_i)$ is constant. Pick any leaf L of Λ_- and any i . There is a directed edge e of G_i that is traversed twice in the same direction by the realization of L in G_i . Surgery of L along e , as in the proof of Lemma 5.1, produces a legal loop in G_i , which maps to a legal loop in each G_j under g_{ji} . Choose $\gamma \in F_r$ representing this loop. The axes of γ in T_i all correspond under the maps \tilde{g}_{ji} , and these axes are all legal, and so $\ell_\gamma(T_i) = \ell_\gamma(T_+)$, completing the proof that \hat{G}_i converges to some $\hat{G} \in \mathcal{X}_r$.

Choose Λ_- isometries $f_i^+: T_i \rightarrow T_+$. We show that, in some sense, a subsequence of f_i^+ converges to a Λ_- isometry $T \mapsto T_+$, proving that $\hat{G} \in \hat{\mathcal{A}}$.

Since G is in the closure of the open cell of \mathcal{X}_r containing all of the G_i , for each i there exists a unique marked homotopy equivalence $p^i: \hat{G}_i \rightarrow \hat{G}$ which collapses each component of a subforest to a vertex, and which is affine on each noncollapsed edge. This subforest is trivial if and only if G is in the same open cell as all of the G_i . Note that $p^j \circ g_{ji} = p^i$.

Choose a finite set A and an indexing $\{e_i^\alpha\}_{\alpha \in A}$ of the edges of \hat{G}_i so that $g_{ji}(e_i^\alpha) = e_j^\alpha$. Let $B \subset A$ correspond to edges that are not collapsed by each p^i . Index the edges of \hat{G} as $\{e^\beta\}_{\beta \in B}$ so that $p^i(e_i^\beta) = e^\beta$. The edge length function of \hat{G}_i , denoted $\ell_i: A \rightarrow (0, +\infty)$, converges to a function $\ell: A \rightarrow [0, \infty)$ which is nonzero precisely on B . By restricting to B , ℓ is identified with the edge length function for \hat{G} .

Choose a marked homotopy equivalence $q^i: \hat{G} \rightarrow \hat{G}_i$ as follows. For each vertex $v \in \hat{G}$, consider the subtree $\tau_i(v) = (p^i)^{-1}(v) \subset \hat{G}_i$, and choose $q^i(v) \in \tau_i(v)$ arbitrarily. For each edge e^β of \hat{G} with endpoints v, w , there is a unique continuous extension of q^i to e^β which is affine on e^β , whose image is contained in $e_i^\beta \cup \tau_i(v) \cup \tau_i(w)$, and such that $p^i \circ q^i|_{e^\beta}$ is homotopic rel endpoints to the identity map. This defines $q^i: \hat{G} \rightarrow \hat{G}_i$ which is clearly a

marked homotopy inverse for p^i . Denote by \tilde{p}^i, \tilde{q}^i the lifts of p^i, q^i to F_r -equivariant maps of universal covers.

Let ϵ_i denote the maximal diameter in \widehat{G}_i of the trees $\tau_i(v)$, and so $\epsilon_i \rightarrow 0$ as $i \rightarrow \infty$. Also let $r_i \geq 1$ denote the maximum value of the ratios

$$\frac{\text{Len } e^\beta}{\text{Len } e_i^\beta}, \frac{\text{Len } e_i^\beta}{\text{Len } e^\beta}, \frac{\text{Len } e^\beta}{\text{Len } e_i^\beta + 2\epsilon_i}, \frac{\text{Len } e_i^\beta + 2\epsilon_i}{\text{Len } e^\beta}$$

as β varies over B , so $r_i \rightarrow 1$ as $i \rightarrow \infty$, and r_i is an upper bound for the amount by which either of the maps p^i, q^i stretches or compresses an edge with index in B ; of course p^i compresses by $+\infty$ any edge with index in $A - B$.

Next we show that for each vertex v of T , the points $f_i^+ \tilde{q}^i(v)$ vary over a compact subset of T_+ . Let $\{E_1, \dots, E_K\}$ be the edges of the tree $\tau_i(v)$. For each E_k choose leaves L_k, L'_k of Λ_- which run through E_k and which are not asymptotic in either direction, and so the realizations of L_k and L'_k in T_+ intersect in a compact interval I_k independent of i . It follows that $f_i^+ \tilde{q}^i(v) \in I_1 \cup \dots \cup I_K$, the latter being compact.

Since there are only finitely many orbits of vertices we can pass to a subsequence so that $f_i^+ \tilde{q}^i(v)$ converges to a point denoted $f(v)$. Extend f affinely to an F_r -equivariant map $f: T \rightarrow T_+$. Since $r_i \rightarrow 1$ it follows that f is an edge isometry. Given a leaf L of Λ_- realized in T , with realizations in T_i, T_+ denoted L_i, L_+ , note that $\tilde{p}^i(L_i) = L$, and that $\tilde{q}^i(L)$ is contained in the ϵ_i neighborhood of L_i . Since $f_i^+ \tilde{q}^i|_{L_i}$ is an isometry onto L_+ , the map $f_i^+ \tilde{q}^i|_L$ has image contained in the ϵ_i neighborhood of L_+ , and so $f(L) = L_+$. The map $f: L \rightarrow L_+$ restricts to an isometry on each edge of T crossed by L , and so this map will be an isometry once we prove that it is injective. The map $f_i^+ \tilde{q}^i|_L$ can only identify a pair of points $x, y \in L$ if there is a vertex v such that $\tilde{q}^i([x, v] \cup [v, y]) \subset \tau_i(v)$, and so

$$\sup_{x, y \in L, f^i \tilde{q}^i(x) = f^i \tilde{q}^i(y)} d_T(x, y) \leq 2r_i \epsilon_i \rightarrow 0 \quad \text{as } i \rightarrow \infty$$

The map $f: L \rightarrow L_+$ is therefore an isometry. Since L is an arbitrary leaf of Λ_- , we have proved that f is a Λ_- isometry, and so $\widehat{G} \in \widehat{\mathcal{A}}$. \diamond

6.2 The Gromov topology on weak train tracks

The proof of Lemma 6.2 item (2) can be interpreted in terms of the Gromov topology on $\widehat{\mathcal{X}}_r$: for any $D \geq 0$ and $\epsilon > 0$, if i is sufficiently large then the restriction of \tilde{q}_i to any subtree of diameter $\leq D$ is an ϵ -almost isometry; we verify this below. In fact, this proves one of the easy directions of Paulin's Theorem [Pau89], namely that if \widehat{G}_i converges to \widehat{G} in the cellular topology then \widehat{G}_i converges to \widehat{G} in the Gromov topology. For use in the proof of Theorem 6.1, in Lemma 6.3 we present a version of this argument tailored specifically to weak train tracks, couched in terms of a relation among weak train tracks called the " Λ_- relation". This result will be applied in the verification of step 3(d) in Section 6.4.

Given a weak train track $T \in \widehat{\mathcal{A}}$, from among the set $I(T, T_+)$ of all Λ_- isometries $T \mapsto T_+$ we pick out one, denoted k_T^+ , as follows. If Λ_- is nonorientable then Theorem 5.8 says that $I(T, T_+)$ has a unique element, taken to be k_T^+ . If Λ_- is orientable then, fixing an orientation, Theorem 5.8 says that $I(T, T_+)$ is a point or a compact oriented arc, and we take k_T^+ to be the positive endpoint of $I(T, T_+)$, characterized by T having a vertex with two or more

positive gates with respect to k_T^+ . Choosing the positive rather than the negative endpoint of $I(T, T_+)$ is arbitrary, the point being to make a choice which is sufficiently canonical to force the appropriate continuity property — namely, Lemma 6.3 — to be true.

Given two weak train tracks $T, T' \in \widehat{\mathcal{A}}$, the Λ_- relation between T and T' is the set $R \subset T \times T'$ consisting of all (x, x') with the following two properties:

- (1) $k_T^+(x) = k_{T'}^+(x')$
- (2) There exists a leaf L of Λ_- such that x is contained in the realization of L in T and x' is contained in the realization of L in T' .

The following lemma says that the Λ_- relation can be used to exhibit closeness among weak train tracks in the Gromov topology:

Lemma 6.3. *For any $T \in \widehat{\mathcal{A}}$, $\epsilon > 0$, any finite subset $t \subset T$, and any finite subset $P \subset F_r$, there exists a neighborhood $U \subset \widehat{\mathcal{A}}$ of T such that for any $T' \in U$, if R denotes the Λ_- relation between T and T' , and if we define a finite subset*

$$t' = \{x' \in T' \mid \text{there exists } x \in t \text{ such that } (x, x') \in R\}$$

then the relation $R \cap (t \times t')$ is a P -equivariant ϵ -almost isometry between t and t' .

Proof. Choose a sequence T_i converging to T . We use various notations, and we know various facts, from the proof of Lemma 6.2 item (2), in particular: marked graphs $\widehat{G}_i = T_i/F_r$, $\widehat{G} = T/F_r$, marked homotopy equivalences $p_i: \widehat{G}_i \rightarrow \widehat{G}$, $q_i: \widehat{G} \rightarrow \widehat{G}_i$ lifting to F_r equivariant maps $\tilde{p}_i: T_i \rightarrow T$, $\tilde{q}_i: T \rightarrow T_i$, etc. Sequences $r_i \geq 1$ converging to 1 and $\epsilon_i \geq 0$ converging to 0 were shown to have the property that for any $\eta > 0$ there exists a neighborhood U of T such that if $T_i \in U$ then $|r_i - 1|, |\epsilon_i| < \eta$. It therefore suffices to prove that if r_i is sufficiently close to 1 and ϵ_i is sufficiently close to 0 then the conclusions of the lemma hold.

Consider now the Λ_- isometries $k_i^+: T_i \rightarrow T_+$. We know that $k_i^+ \tilde{q}_i: T \rightarrow T_+$ converges to a Λ_- isometry $f_T^+: T \rightarrow T_+$. We wish to show that f_T^+ has a vertex with two or more positive gates, allowing us to conclude that $f_T^+ = k_T^+$. Since there are only finitely many orbits of turns under the action of F_r , we may pass to a subsequence so that the gates in T_i with respect to the maps k_i^+ all correspond under the F_r -equivariant homeomorphisms \tilde{g}_{ji} . Pick a vertex $v \in T$ so that each $\tau_i(v)$ contains a vertex w_i with two or more positive gates relative to k_i^+ . The construction of the maps \tilde{q}_i gives us the freedom to choose $\tilde{q}_i(v)$ in $\tau_i(v)$, and we choose $\tilde{q}_i(v) = w_i$. It follows that there is a nondegenerate turn $\{e, e'\}$ at v such that the directions $D_{\tilde{q}_i}(e)$ and $D_{\tilde{q}_i}(e')$ are contained in distinct gates of k_i^+ at w_i . Passing to the limit, this proves that e, e' are contained in distinct gates of f_T^+ at v , proving that $f_T^+ = k_T^+$.

Let R_i denote the Λ_- relation between T and T_i : $x \in T$ and $y \in T_i$ are related by R if $k_T^+(x) = k_i^+(y)$ and x, y are contained in realizations of the same leaf of Λ_- in T, T_i , respectively. Assuming this is so, we argue that the distance between y and $\tilde{q}^i(x)$ converges to zero uniformly as $i \rightarrow +\infty$, independent of the choice of $(x, y) \in R$. Pick any leaf L of Λ_- realized in T , with realization $L_i \subset T_i$ and $L_+ \subset T_+$, so that $x \in L$ and $y \in L_i$. The arguments of Lemma 6.2 item (2) show that $\tilde{q}^i(x)$ is within distance ϵ_i of some point $z_i \in L_i$, and so $k_i^+ \tilde{q}^i(x)$ is within distance ϵ_i of $k_i^+(z_i) \in L_+$. On the other hand $k_i^+ \tilde{q}^i(x)$ converges uniformly to $k_T^+(x) = k_i^+(y) \in L_+$, and so $k_i^+(z_i)$ converges to $k_i^+(y)$ in L_+ , and thus the distance between $k_i^+(z_i)$ and $k_i^+(y)$ converges uniformly to zero. Since k_i^+ restricts to an

isometry from L_i to L_+ , the distance between y and z_i in L_i converges uniformly to zero. This proves that the distance between $\tilde{q}^i(x)$ and y converges uniformly to zero, as desired.

Pick $\eta > 0$, pick a finite subset $t \subset T$, and let $t_i \subset T_i$ be the finite set of points $y \in T_i$ related to some $x \in t$ by the relation R_i . The set $R_i \cap (t \times t_i)$ projects surjectively to t and to T_i because $R_i \subset T \times T_i$ projects surjectively to T and to T_i . Let D be the diameter of t . We must show that for sufficiently large i the relation $R_i \cap (t \times t_i)$ is an η -almost isometry (equivariance with respect to any finite subset of F_r is easily verified), in other words if $(x, y), (x', y') \in R_i \cap (t \times t_i)$ then

$$|d_T(x, x') - d_{T_i}(x', y')| < \eta$$

Picking i sufficiently large so that $d_{T_i}(\tilde{q}^i(x), y) < \frac{\eta}{4}$ for any $(x, y) \in R_i$, it remains to show that the restriction $\tilde{q}^i \mid t$ is an $\frac{\eta}{2}$ -almost isometry, that is, if $x, x' \in t$ then

$$d_T(x, x') - \frac{\eta}{2} \leq d_{T_i}(\tilde{q}^i(x), \tilde{q}^i(x')) \leq d_T(x, x') + \frac{\eta}{2} \quad (6.1)$$

There is an integer $J \geq 0$ such that for any $x, x' \in t$ the arc $[x, x']$ makes at most J turns. The restriction $\tilde{q}^i \mid [x, x']$ stretches or compresses each edge by at most $r_i \geq 1$, and makes a fold at each turn of length at most $\epsilon_i \geq 0$, as shown in the proof of Lemma 6.2 item 2. It follows that

$$\frac{1}{r_i} d_T(x, x') - J\epsilon_i \leq d_{T_i}(\tilde{q}^i(x), \tilde{q}^i(x')) \leq r_i d_T(x, x')$$

The right hand inequality of (6.1) follows by choosing i so large that $r_i < 1 + \frac{\eta}{2D}$, from which it follows that $r_i d_T(x, x') \leq d_T(x, x') + \frac{\eta}{2}$. The left hand inequality follows by first choosing i so large that $\frac{\eta}{2} - J\epsilon_i > 0$, and then choosing i so large that $(1 - \frac{1}{r_i})D \leq \frac{\eta}{2} - J\epsilon_i$ which implies that $d_T(x, x') - \frac{\eta}{2} \leq \frac{1}{r_i} d_T(x, x') - J\epsilon_i$. \diamond

Remarks. In the course of this proof we have shown a little more regarding the structure of the space $I(\hat{\mathcal{A}}, T_+) = \cup_{T \in \hat{\mathcal{A}}} I(T, T_+)$ in the Gromov topology (we refer the reader to [Sko], or to Section 6.4, verification of Step 3(d), for the definition of the Gromov topology on such spaces of maps). There is a continuous, open, surjective map $\pi: I(\hat{\mathcal{A}}, T_+) \rightarrow \hat{\mathcal{A}}$, the “domain map”, which maps each function $f: T \rightarrow T_+$ in $I(\hat{\mathcal{A}}, T_+)$ to its domain T in $\hat{\mathcal{A}}$. In the above proof, we have selected a particular element k_T^+ from each fiber $\pi^{-1}(T) = I(T, T_+)$, thereby defining a section of the map π . What we have proved amounts to the statement that this section is a continuous map $k^+: \hat{\mathcal{A}} \rightarrow I(\hat{\mathcal{A}}, T_+)$. In the case that Λ_- is nonorientable one has in fact that π is a homeomorphism with inverse k^+ . But when Λ_- is orientable and when there exist nontrivial fibers, one can also select the negative endpoint, and the same proof gives a continuous section $k^-: \hat{\mathcal{A}} \rightarrow I(\hat{\mathcal{A}}, T_+)$. This shows that, in some sense, the intervals $I(T, T_+)$ all fit together in a continuous fashion, making the space $I(\hat{\mathcal{A}}, T_+)$ into an “interval bundle” over $\hat{\mathcal{A}}$ whose fiber over $T \in \hat{\mathcal{A}}$ is $I(T, T_+)$. To make sense out of this one would have to investigate local triviality, and one would have to deal with the fact that the fibers may not be topologically homogeneous, some being nondegenerate intervals and others degenerating to points. We shall not elaborate on this further, except insofar as to point out that the space of efficient representatives studied by Los and Lustig in [LL04] is naturally a subspace of the total space $I(\hat{\mathcal{A}}, T_+)$ of this interval bundle.

6.3 Properness of the length map.

Our present goal, a piece of the proof that $\text{Len}: \mathcal{A} \rightarrow (0, +\infty)$ is a proper homotopy equivalence, is the following:

Proposition 6.4 (Length is proper). *The map $\text{Len}: \mathcal{A} \rightarrow (0, +\infty)$ is proper. It follows that a subset $C \subset \mathcal{A}$ is compact if and only if it is closed in \mathcal{A} and the set $\text{Len}(C)$ is contained in a compact subinterval of $(0, +\infty)$.*

The next lemma gives a compactness criterion that will be used several times in the proof of Proposition 6.4. Recall our convention from Section 2.2 that all fold maps are edge isometries. Given a subset $K \subset \widehat{\mathcal{X}}_r$ define $F(K; 1)$ to be the union of K with the set of all $H \in \widehat{\mathcal{X}}_r$ for which there exists $H_0 \in K$ and a fold map $f: H_0 \rightarrow H$. By induction define $F(K; M) = F(F(K; M-1); 1)$; equivalently this is the set of all $H \in \widehat{\mathcal{X}}_r$ for which there exists $H_0 \in K$ and a piecewise isometry $f: H_0 \rightarrow H$ (not required to be an edge isometry) which factors into M marking preserving homotopy equivalences $H_0 \mapsto H_1 \mapsto \dots \mapsto H_M = H$ each of which is either an isometry or a fold map.

Lemma 6.5. *For each compact $K \subset \widehat{\mathcal{X}}_r$ and integer $M \geq 0$ the set $F(K, M)$ is compact.*

Proof. By induction it suffices to prove that $F(K; 1)$ is compact.

Let each closed cell σ of \mathcal{X}_r be labelled $\sigma(G)$ for some choice of marked graph $G \in \text{int}(\widehat{\sigma})$, so if $H \in \mathcal{X}_r$ then $H \in \text{int}(\sigma(G))$ if and only if there exists a marking preserving homeomorphism $G \approx H$; moreover, such a homeomorphism is unique up to isotopy rel vertices. Let $\widehat{\sigma}(G)$ denote the preimage of $\sigma(G)$ in $\widehat{\mathcal{X}}_r$. By local finiteness of the cell decomposition of \mathcal{X}_r , the compact set K intersects only finitely many closed cells of $\widehat{\mathcal{X}}_r$, and by intersecting each of these one at a time it suffices to fix a closed cell $\widehat{\sigma}(G)$ and to assume that $K \subset \widehat{\sigma}(G)$.

For each ordered turn (e, e') of G let $F(K; e, e')$ denote the union of K with the set of all $H \in \widehat{\mathcal{X}}_r$ for which there exists $H_0 \in K$ and a fold map $f: H_0 \mapsto H$ such that, with respect to the marked homeomorphism $G \approx H_0$, the map f folds the turn (e, e') and the length of e in H_0 is greater than or equal to the length of e' in H_0 . By finiteness of the number of turns of G it suffices to fix an ordered turn (e, e') of G and to prove that $F(K; e, e')$ is compact.

We consider two cases, depending on the valence of the common initial vertex v of the edges e, e' .

Suppose first that v has valence ≥ 4 . Let $G^\#$ be the marked graph obtained from G by folding proper initial segments of e and of e' to obtain an edge $e^\#$ of $G^\#$. The cell $\widehat{\sigma}(G)$ is a codimension 1 face of $\widehat{\sigma}(G^\#)$, defined by a marked homotopy equivalence $G^\# \rightarrow G$ that collapses $e^\#$ to a point and is otherwise a bijection. Under this collapse map, the edges e, e' of G correspond to two edges of $G^\#$ also denoted e, e' . Let $\Delta = \{(H_0, t) \in \widehat{\sigma}(G) \times [0, \infty) \mid 0 \leq t \leq \text{Len}_{H_0}(e') \leq \text{Len}_{H_0}(e)\}$. Define a map $\Delta \mapsto \widehat{\sigma}(G^\#)$ taking (H_0, t) to the element of $\widehat{\sigma}(G^\#)$ which assigns length t to $e^\#$, which subtracts t from the lengths of the edges e, e' in H_0 , and which leaves the lengths of all other edges of H_0 unchanged. Notice that when $t = 0$ the length assigned to $e^\#$ equals zero; whereas when $t = \text{Len}_{H_0}(e')$ then the length assigned to e' equals zero and, in the case $\text{Len}_{H_0}(e') = \text{Len}_{H_0}(e)$, the length assigned to e also equals zero; and so in either of these two cases the image of (H_0, t) is in a face of the closed cell $\widehat{\sigma}(G^\#)$. The map $\Delta \mapsto \widehat{\sigma}(G^\#)$ is clearly continuous, and $F(K; e, e')$ is the image under this map of a compact subset of Δ , namely the set $\{(H_0, t) \in \Delta \mid H_0 \in K\}$, proving that $F(K; e, e')$ is compact.

Suppose next that v has valence 3. Let e'' be the third edge incident to v . Note that e'' is distinct from both e and e' , for if say $e'' = e$, consider any oriented leaf L of Λ_- travelling backwards along e ; when L hits the vertex v , it cannot take the illegal turn (e, e') and so L must continue into $e'' = e$ again, and this must repeat forever, contradicting surjectivity of half-leaves of Λ_- . Now define Δ as above, and define the map $\Delta \mapsto \hat{\sigma}(G)$ taking (H_0, t) to the element of $\hat{\sigma}(G)$ which subtracts t from the lengths of e, e' and adds t to the length of e'' . The proof is then completed as in the previous case. \diamond

Proof of Proposition 6.4. Since $\hat{\mathcal{A}}$ is a closed subset of $\hat{\mathcal{X}}_r$ and since Len is continuous, for each closed interval $[a, b] \subset (0, +\infty)$ the set $\hat{\mathcal{A}}_{[a, b]} = \hat{\mathcal{A}} \cap \text{Len}^{-1}[a, b]$ is closed, and it suffices to prove that $\hat{\mathcal{A}}_{[a, b]}$ is compact. Lemma 6.6 below implies that there exists a train track $G \in \hat{\mathcal{A}}$ and an integer $M \geq 0$ such that for each $H \in \hat{\mathcal{A}}_{[a, b]}$ there is a Λ_- isometry $G \mapsto H$ which has a Stallings fold sequence consisting of at most M folds. It follows that $\hat{\mathcal{A}}_{[a, b]} \subset F(\{G\}; M)$, and the latter set is compact by Lemma 6.5. Since $\hat{\mathcal{A}}_{[a, b]}$ is a closed subset of a compact set, it is compact. This completes the proof subject to our verifying Lemma 6.6. \diamond

Lemma 6.6. *For any closed interval $[a, b] \subset (0, +\infty)$ there exists a train track map $g: G \rightarrow G$ with G normalized to be in $\hat{\mathcal{A}}$, and there exists an integer $M \geq 0$, so that if $H \in \hat{\mathcal{A}}_{[a, b]}$, and if $f_{\tilde{H}}: \tilde{H} \rightarrow T_+$ is a Λ_- -isometry, then there exists an edge isometry $h: G \rightarrow H$ such that $f_g = f_{\tilde{H}}h$, and for any such edge isometry $h: G \rightarrow H$, the combinatorial length of h is at most M , and h has a Stallings fold sequence of length at most M .*

Proof. Pick a train track $G \in \mathcal{A}$ with the finest possible Nielsen structure. Pick $\epsilon > 0$ for G by applying Proposition 5.4. We may assume that $\epsilon > b$, for otherwise we replace G by $\phi^{-n}(G)$ for sufficiently large n , which has the effect of replacing ϵ by $\epsilon \cdot \lambda(\phi)^n$. By Proposition 5.4, for every $H \in \mathcal{A}_{[a, b]}$ and every Λ_- isometry $f_{\tilde{H}}: \tilde{H} \rightarrow T_+$ there is an edge isometry $h: G \rightarrow H$ so that $f_g = f_{\tilde{H}}h$ where \tilde{h} is a lift of $h: G \rightarrow H$.

Let $h: G \rightarrow H$ be any edge isometry such that $f_g = f_{\tilde{H}}h$. Recall from the discussion preceding Lemma 2.6 that, after subdividing H at the h -images of the vertices of G , the subdivided H has at most $5r - 5$ edges. Since $\text{Len}(H) \geq a$, there exists an edge e of H with $\text{Len}(e) \geq a/(5r - 5)$. For any edge isometry $G' \rightarrow H$ whose vertices map disjointly from e , define an e -arc in G' to be an open arc of G' which maps isometrically onto the interior of e . Two e -arcs in G' are disjoint and each e -arc has length equal to $\text{Len}(e)$, and so G' has at most $\text{Len}(G')(5r - 5)/a$ distinct e -arcs.

Using the notation of (2.2), apply Lemma 2.6 to obtain a Stallings fold sequence $G = G_0 \xrightarrow{g_1} G_1 \xrightarrow{g_2} \cdots \xrightarrow{g_J} G_J = H$ for h whose length J is bounded by the combinatorial length of h . For $j = 0, \dots, J$ let M_j be the number of e -arcs in G_j . The sequence M_j is nonincreasing: each e -arc in G_{j-1} maps isometrically onto an e -arc in G_j ; the e -arcs in G_{j-1} are precisely the components of the inverse images of the e -arcs of G_j ; and so either $M_j = M_{j-1}$, or $M_j < M_{j-1}$ which occurs precisely when the fold map $g_j: G_{j-1} \rightarrow G_j$ identifies at least one pair of e -arcs in G_{j-1} to a single e -arc of G_j . There is therefore a sequence $0 \leq j_1 < \cdots < j_N = J$ of some length N such that the subsequence M_j , for $j_{n-1} < j \leq j_n$, is constant, and such that $M_{j_n} > M_{j_{n+1}}$. It follows that $N \leq M_0 \leq \text{Len}(G)(5r - 5)/a$.

The following lemma, applied with K equal to the singleton $\{G\}$, and with h equal to the composed map $G_0 \mapsto G_{j_1-1}$, implies that G_{j_1-1} belongs to a compact subset K'

of $\widehat{\mathcal{X}}_r$ and that the map $G_0 \mapsto G_{j_1-1}$ has uniformly bounded combinatorial length. Then apply Lemma 6.5 to the single fold $G_{j_1-1} \mapsto G_{j_1}$ to obtain a compact set K'' that contains G_{j_1} . Then reapply the following lemma with $K = K''$ and with h equal to the composed map $G_{j_1} \mapsto G_{j_2-1}$. Continue inductively in this manner to prove that each composed map $G_{j_{n-1}} \mapsto G_{j_n-1}$ has uniformly bounded combinatorial length. Summing up we obtain a uniform bound M for the combinatorial length of h . Applying Lemma 2.6, h has a Stallings fold sequence of length at most M . \diamond

Lemma 6.7. *For any compact set $K \subset \widehat{\mathcal{X}}_r$ there exists a compact set $K' \subset \widehat{\mathcal{X}}_r$ and an integer $L \geq 0$ so that for any $G \in K$, $G' \in \widehat{\mathcal{X}}_r$, and edge isometry $h: G \rightarrow G'$, if $\ell_\gamma(G') \geq \ell_\gamma(T^+)$ for all $\gamma \in F_r$, and if there exists an edge of G' over the interior of which h is injective, then $G' \in K'$ and h has combinatorial length $\leq L$.*

In this lemma, edges of G' are determined as usual by subdividing at the images of the vertices of G under the map h , and so G' has at most $5r - 5$ edges.

Proof. A proper, connected, nontree subgraph of a marked graph G determines a conjugacy class of proper, nontrivial free factors of F_r by pulling back the fundamental group of the subgraph under the marking $R_r \rightarrow G$. Let $\mathcal{FF}(K)$ be the set of all free factor conjugacy classes associated to the proper, connected, nontree subgraphs of G , as G varies over K . This set is finite, because the cells of the cellular topology on $\widehat{\mathcal{X}}_r$ are locally finite (see Section 2.1). By Theorem 5.4 of [BFH97], for each proper free factor conjugacy class A in F_r , the infimum of $\ell_\gamma(T^+)$ for $\gamma \in A$ is a positive number denoted $\epsilon(A)$. Let $\epsilon > 0$ be the minimum value of $\epsilon(A)$ as A varies over the finitely many elements of $\mathcal{FF}(K)$. If G, G' are as in the lemma then it follows that for each $\gamma \in F_r$, if the circuit in G representing γ does not surject onto G , then the circuit in G' representing γ has length $\geq \epsilon$.

Let $e' \subset G'$ be an edge over whose interior h is injective, let X' be the complement in G' of the interior of e' , and let $X = h^{-1}(X') \subset G$. Then $h(X) = X'$ and, since h is a homotopy equivalence, h restricts to a homotopy equivalence between X and X' . It follows that X, X' have the same number of components — two or one depending on whether e' is separating or nonseparating — these components are connected, proper, subgraphs of G, G' , respectively, and neither X nor X' is a tree for otherwise G or G' would have a valence 1 vertex. Moreover, components of X, X' that correspond under h represent the same proper, nontrivial free factor conjugacy class of F_r . Thus every simple circuit in X' has length at least ϵ and so crosses an edge of length at least $\epsilon/(5r - 6)$, the subgraph X' having at most $5r - 6$ edges. Letting E' be the union of e' and the set of edges of G' with length at least $\epsilon/(5r - 6)$, it follows that every simple circuit in G' intersects an edge in E' , and so the complement of E' is a forest U' . For each edge e of G , since $\text{Len}(e) \leq \text{Len}(G)$, the edge path $h(e)$ contains at most $1 + \text{Len}(G) \cdot (5r - 6)/\epsilon$ elements of E' . Removing these edges from the edge path $h(e)$ leaves at most $2 + \text{Len}(G) \cdot (5r - 6)/\epsilon$ complementary subpaths each of which is contained in the forest U' and so is embedded, crossing each edge of U' at most once. We conclude that there is a uniform bound to the combinatorial length of the edge path $h(e)$. Since e was arbitrary, the combinatorial length of h is uniformly bounded. The map h therefore has a Stallings fold sequence of uniformly bounded length. Applying Lemma 6.5, G' lies in a compact subset K' of \mathcal{X}_r . \diamond

6.4 Applying Skora's method to the Properness Theorem 6.1

Let $\mathcal{A} \subset \mathcal{X}_r$ be the set of weak train tracks of a nongeometric, fully irreducible $\phi \in \text{Out}(F_r)$, with a normalization $\mathcal{A} \hookrightarrow \widehat{\mathcal{X}}_r$ determined as usual by a choice of normalization of T_ϕ^+ . Rather than continue with the notation $\widehat{\mathcal{A}}$ for the image of the normalization map, in this section we assume implicitly that all marked graphs and F_r -trees representing elements of \mathcal{A} are normalized.

The major tool that we use in proving Theorem 6.1, that the map $\text{Len}: \mathcal{A} \mapsto (0, +\infty)$ is a proper homotopy equivalence, is Skora's method [Sko] for detecting the homotopy type of various spaces of trees. The preprint [Sko] being unpublished and therefore difficult to obtain, other accounts to consult include [Whi93], [Cla05], and [GL05]; we shall quote a technical result from the latter which simplifies the method. While Skora applies his technique to very general spaces of group actions on trees, we shall consider only free, simplicial, minimal F_r -trees, that is, elements of outer space \mathcal{X}_r .

Spaces of edge isometries. Skora's method requires defining topologies on various spaces of piecewise isometries between trees. We adopt here an addition to the method suggested by Definition 3.8 and Proposition 3.10 of [GL05] which guarantees that minimality of trees is true at all stages of Skora's method, and which allows us to focus solely on edge isometries among free, simplicial, minimal F_r trees, that is, elements of $\widehat{\mathcal{X}}_r$.

Given free, simplicial, minimal F_r trees S, T , a piecewise isometry $f: S \rightarrow T$ satisfies the *minimality condition* if for each $x \in S$ there exists an open arc embedded in S and containing x such that the restriction of f to this arc is an isometry; equivalently, f is an edge isometry and each vertex of S has at least two gates with respect to S . The motivation for the minimality condition is that if it failed, if there were a point x having a single gate, then one could construct an equivariant Stallings fold sequence for f by first folding all of the turns at x , producing a tree which has a valence 1 vertex and therefore is not minimal. Particularly important is the fact that Λ_- isometries automatically satisfy the minimality condition.

Let \mathcal{M} be the set of equivalence classes of piecewise isometries $f: S \rightarrow T$ satisfying the minimality condition, where S, T are free, simplicial, minimal F_r trees, and where $f: S \rightarrow T$ and $f': S' \rightarrow T'$ are equivalent if there exist F_r equivariant isometries $h: S \rightarrow S', k: T \rightarrow T'$ such that $kf = f'h$. Formally the element $f: S \rightarrow T$ is to be regarded as an ordered triple (f, S, T) which allows us to define two projections $\text{Dom}, \text{Ran}: \mathcal{M} \rightarrow \mathcal{X}_r$, the domain projection $\text{Dom}(f, S, T) = S$ and the range projection $\text{Ran}(f, S, T) \mapsto T$.

The group $\text{Out}(F_r)$ acts on \mathcal{M} from the right so that the domain and range projections are $\text{Out}(F_r)$ equivariant. The action is defined as follows. Given $f: S \rightarrow T$ in \mathcal{M} and given $\phi \in \text{Out}(F_r)$, choose $\Phi \in \text{Aut}(F_r)$ representing ϕ . Recall that the F_r -trees $S \cdot \phi, T \cdot \phi$ have underlying \mathbf{R} -trees S, T but with F_r actions precomposed by Φ . The F_r -equivariant map $f: S \rightarrow T$ therefore remains F_r -equivariant when the actions are precomposed by Φ , and we define this map to be the image of f under ϕ , well-defined in \mathcal{M} independent of the choice of Φ .

We put off until later the definition of the Gromov topology on \mathcal{M} , as given in the references above. For the moment, we note that the domain and range projections $\text{Dom}, \text{Ran}: \mathcal{M} \rightarrow \mathcal{X}_r$ are continuous maps with respect to this topology.

Notice that the reparameterized maps $-\log \text{Len} = \log\left(\frac{1}{\text{Len}}\right) : \mathcal{A} \rightarrow \mathbf{R}$ and $\frac{1}{\text{Len}} : \mathcal{A} \rightarrow (0, +\infty)$ are “correctly oriented” in the sense that, on an orbit $G_i = \phi^i(G_0)$, the values of these maps increase as i increases. The map $\text{Len} : \mathcal{A} \rightarrow (0, +\infty)$, on the other hand, is “incorrectly oriented”, and the map $-\log \text{Len}$ carries heavier notational baggage, so in the proof that follows we often think in terms of the map $\frac{1}{\text{Len}}$.

Proof of the Properness Theorem 6.1. We give the proof in the several steps.

Step 1: A bi-infinite fold path \mathcal{L} . Our goal being to prove that $\frac{1}{\text{Len}} : \mathcal{A} \rightarrow (0, +\infty)$ is a proper homotopy equivalence, we begin by choosing a subset $\mathcal{L} \subset \mathcal{A}$ such that $\frac{1}{\text{Len}} : \mathcal{L} \rightarrow (0, +\infty)$ is a homeomorphism. This subset \mathcal{L} will be a particular fold line, periodic under an appropriate power of ϕ .

Passing to a positive power of ϕ , we may assume ϕ is rotationless. Applying Proposition 4.4, choose $g : G \rightarrow G$ to be a train track representative of ϕ whose local Whitehead decomposition is as fine as possible, meaning that the local stable Whitehead graphs of g have no cut points. We assume that the train track G is normalized just as \mathcal{A} is, and so there is an associated Λ_- isometry $f_g : \tilde{G} \rightarrow T_+$. Define $G_0 = G$ and $G_i = \lambda^{-i} \phi^i(G)$ for all $i \in \mathbf{Z}$. The map $g : G \rightarrow G$ induces marked edge isometries $g_{i+1,i} : G_i \rightarrow G_{i+1}$. Let $L_i = \text{Len}(G_i) = \lambda^{-i} \text{Len}(G_0)$.

For later use recall that if we define $g_{ji} = g_{j,j-1} \circ g_{j-1,j-2} \circ \cdots \circ g_{i+1,i} : G_i \rightarrow G_j$ then the universal covering trees \tilde{G}_i and the F_r -equivariant edge isometries \tilde{g}_{ji} form a direct system in the category of metric spaces and distance nonincreasing maps, the direct limit of this system is the F_r tree T_+ , we obtain direct limit maps $\tilde{g}_i^+ : \tilde{G}_i \rightarrow T_+$ such that the two maps $f_g, \tilde{g}_0^+ : \tilde{G} = \tilde{G}_0 \rightarrow T_+$ are equal to each other.

As described in Section 2.2, choose a Stallings fold sequence for $g_{1,0}$ and interpolate the fold maps to get a Stallings fold path for $g_{1,0}$ connecting G_0 to G_1 . We parameterize this fold path by $\frac{1}{\text{Len}}$, with parameter interval $[\frac{1}{L_0}, \frac{1}{L_1} = \frac{\lambda}{L_0}]$. Translating this path by ϕ^i and reparameterizing, one obtains a Stallings fold path for $g_{i-1,i}$ connecting G_{i-1} to G_i and parameterized by $\frac{1}{\text{Len}}$, with parameter interval $[\frac{\lambda^i}{L_0}, \frac{\lambda^{i+1}}{L_0}]$. Concatenating these paths over all $i \in \mathbf{Z}$ we obtain a bi-infinite Stallings fold path $(0, +\infty) \mapsto \mathcal{A}$. The image of this path is the desired set \mathcal{L} , completing Step 1.

Note that $\frac{1}{\text{Len}} : \mathcal{L} \rightarrow (0, +\infty)$ is a homeomorphism. The inverse of this homeomorphism, which we denote $\rho : (0, +\infty) \rightarrow \mathcal{L}$, will be our preferred parameterization for the fold path \mathcal{L} .

Step 2: The map $\mathcal{A} \rightarrow \mathcal{L}$. We construct a proper, continuous map $\sigma : \mathcal{A} \rightarrow \mathcal{L}$ whose restriction to \mathcal{L} , while not the identity on \mathcal{L} , is properly homotopic to the identity on \mathcal{L} . In later steps we shall prove that σ is a proper homotopy equivalence.

Applying Proposition 5.4, there exists a number $\eta > 0$ such that if T is any weak train track with $\text{Len}(T) < \eta$ and if $f_T : T \rightarrow T_+$ is a Λ_- isometry, then there exists a Λ_- isometry $\psi : \tilde{G}_0 \rightarrow T$ such that $f_T \circ \psi = f_g$.

Define the map $\sigma : \mathcal{A} \rightarrow \mathcal{L}$ by

$$\sigma(T) = \rho\left(\frac{\eta}{L_0 \lambda \text{Len}(T)}\right) = \rho\left(\frac{\eta}{L_0 \lambda} \rho^{-1}(T)\right)$$

This map is evidently continuous, and it is proper by Proposition 6.4. The restriction to \mathcal{L} , when conjugated by ρ , is the map from $(0, +\infty)$ to itself given by $r \mapsto \frac{\eta}{L_0 \lambda} r$. This map being an orientation preserving homeomorphism of $(0, +\infty)$, the map $\sigma \mid \mathcal{L}: \mathcal{L} \rightarrow \mathcal{L}$ is properly homotopic to the identity.

Remark. For later purposes we note that $\eta < L_0 < \lambda L_0$ and so $\frac{\eta}{\lambda L_0} < 1$. In other words, the map σ decreases the value of the parameter $\frac{1}{\text{Len}}$, multiplying it by the number $\frac{\eta}{\lambda L_0}$, and so σ moves each point of \mathcal{L} strictly to the left.

Step 3: The map to \mathcal{M} . For each $T \in \mathcal{A}$ we define a Λ_- isometry $\psi^T: \sigma(T) \rightarrow T$ with the property that the map $(\sigma(T), T) \mapsto \psi^T$ is a continuous map from the set $\{(\sigma(T), T) \in \mathcal{L} \times \mathcal{A} \mid T \in \mathcal{A}\}$ to the set \mathcal{M} . Letting $u \in (0, +\infty)$ be such that $\sigma(T) = \rho(u)$, the map $\psi^T: \sigma(T) \rightarrow T$ will be defined by factoring a particular Λ_- isometry $f_u^+: \rho(u) = \sigma(T) \rightarrow T_+$ through T , where the other factor is a particular Λ_- isometry $k_T^+: T \rightarrow T_+$. We construct these Λ_- isometries, and we define ψ^T and verify its properties, in several substeps.

Step 3(a): The Λ_- isometry f_u^+ from $\rho(u) \in \mathcal{L}$ to T_+ . Define the *semiflow* of the flow line \mathcal{L} as follows. This semiflow will be a family of F_r -equivariant Λ_- isometries $f_{ts}: \rho(t) \rightarrow \rho(s)$, $0 < s < t < +\infty$, so that the semiflow identity $f_{ut} \circ f_{ts} = f_{us}$ is satisfied. For each integer i the map $f_{1/L_i, 1/L_{i-1}}: \rho(1/L_{i-1}) \rightarrow \rho(1/L_i)$ is defined to be the map $\tilde{g}_{i,i-1}: \tilde{G}_{i-1} \rightarrow \tilde{G}_i$. For each integer i , the subfamily of maps f_{ts} for $\frac{1}{L_{i-1}} \leq s \leq t \leq \frac{1}{L_i}$ is (up to reparameterization) defined to be the semiflow of the chosen Stallings fold sequence for $g_{i,i-1}: G_i \rightarrow G_{i-1}$. The definition of f_{ts} extends uniquely to $0 < s \leq t < \infty$ in such a way as to satisfy the semiflow identity.

The trees $\rho(t)$ and the maps f_{ts} form a direct system in the category of F_r -equivariant metric spaces and distance nonincreasing maps. The direct limit of this direct system is the F_r tree T_+ , because the trees \tilde{G}_i and the maps g_{ji} form a co-finite subsystem, whose direct limit is known to be T_+ by Theorem 2.13. We obtain direct limit maps $f_u^+: \rho(u) \rightarrow T_+$ which satisfy the property that $f_u^+ \circ f_{ut} = f_t^+$ for $0 < t < u < +\infty$. When $\rho(u) = \sigma(T)$ we also denote $f_u^+ = f_{\sigma(T)}^+: \sigma(T) \rightarrow T_+$.

Step 3(b): The Λ_- isometry k_T^+ from $T \in \mathcal{A}$ to T_+ . This was defined in Section 6.2 using Theorem 5.8: if Λ_- is nonorientable then k_T^+ is the unique Λ_- isometry $T \mapsto T_+$; and when Λ_- is oriented then k_T^+ is the positive endpoint of the compact oriented interval $I(T, T_+)$ of Λ_- isometries.

Step 3(c): The Λ_- isometry $\psi^T: \sigma(T) \rightarrow T$. For each $T \in \mathcal{A}$ the map ψ^T is characterized as the unique Λ_- isometry from $\sigma(T)$ to T with the property that $f_{\sigma(T)}^+ = k_T^+ \circ \psi^T$. Uniqueness of ψ^T follows from the proof of uniqueness in Proposition 5.4. However, we still have to establish that ψ^T exists with this property.

The existence of ψ^T being invariant under the action of ϕ on \mathcal{A} , it suffices to prove existence after replacing T by $\phi^k(T)$ for a particular power k . Since ϕ increases the parameter

$\frac{1}{\text{Len}}$ by multiplying it by the quantity λ , after acting by the appropriate power of ϕ , we may assume that

$$\frac{1}{\eta} \leq \frac{1}{\text{Len}(T)} \leq \frac{\lambda}{\eta}$$

Recalling that $\frac{1}{\text{Len}(\sigma(T))} = \frac{\eta}{L_0 \lambda \text{Len}(T)}$ we have

$$\frac{1}{L_0 \lambda} \leq \frac{1}{\text{Len}(\sigma(T))} \leq \frac{1}{L_0}$$

Setting $t_0 = \frac{1}{\text{Len}(\sigma(T))}$, the semiflow map $f_{\frac{1}{L_0}, t_0}: \sigma(T) \rightarrow \tilde{G}$ is therefore defined. Since $\text{Len}(T) \leq \eta$, Proposition 5.4 gives us a Λ_- isometry $h_T: \tilde{G} = \rho(1/L_0) \rightarrow T$ with the property that $f_g = k_T^+ \circ h_T$. Define $\psi^T = h_T \circ f_{\frac{1}{L_0}, t_0}$, and so

$$\begin{aligned} k_T^+ \circ \psi^T &= k_T^+ \circ h_T \circ f_{\frac{1}{L_0}, t_0} \\ &= f_g \circ f_{\frac{1}{L_0}, t_0} \\ &= f_{\frac{1}{L_0}}^+ \circ f_{\frac{1}{L_0}, t_0} \\ &= f_{t_0}^+ = f_{\sigma(T)}^+ \end{aligned}$$

as desired.

Remark: What we have done is to define ψ^T first for T in a fundamental domain of the action of ϕ on \mathcal{A} , and then to extend over all of \mathcal{A} by ϕ -equivariance. This raises the issue of well-definedness of ψ^T where one end of the fundamental domain, defined by $\frac{1}{\text{Len}(T)} = \frac{1}{\eta}$, is translated by ϕ to the other end of the fundamental domain, defined by $\frac{1}{\text{Len}(T)} = \frac{\lambda}{\eta}$. This issue is made moot by the fact that ψ^T is already uniquely characterized: it is the unique Λ_- isometry from $\sigma(T) \rightarrow T$ that satisfies the equation $f_t^+ = k_T^+ \circ \psi^T$.

Step 3(d): Continuity of the map $(\sigma(T), T) \mapsto \psi^T$. We must prove that the map from the set $\{(\sigma(T), T) \mid T \in \mathcal{A}\} \subset \mathcal{L} \times \mathcal{A}$ to the set \mathcal{M} , taking the pair $(\sigma(T), T)$ to the map $\psi^T: \sigma(T) \rightarrow T$, is a continuous map.

This is the most technical of the steps and requires us to look carefully at the topology of the mapping space \mathcal{M} . We put this issue off for a while more, in order to get through the steps in the proof of Theorem 6.1.

Step 4: Combining Steps 2 and 3, we obtain a continuous map $\mathcal{A} \mapsto \mathcal{M}$ defined by $T \mapsto \psi^T$, so that postcomposing with $\text{Dom}: \mathcal{M} \rightarrow \mathcal{X}_r$ we obtain the map $\sigma: \mathcal{A} \mapsto \mathcal{L}$, and postcomposing with the map $\text{Ran}: \mathcal{M} \rightarrow \mathcal{X}_r$ we obtain the identity map on \mathcal{A} .

Now apply the following theorem:

Theorem 6.8 ([Sko], [GL05]). *Letting $\nabla = \{(s, t) \in [0, \infty] \times [0, \infty] \mid s \leq t\}$, there exists a continuous map $\theta: \mathcal{M} \times \nabla \rightarrow \mathcal{M}$, denoted formally as*

$$\theta(\psi, s, t) = (\psi_{ts}, T_s, T_t)$$

with the following properties for each $\psi: S \rightarrow T$ in \mathcal{M} :

- (1) The map $\psi_{\infty,0}: T_0 \rightarrow T_\infty$ equals $\psi: S \rightarrow T$.
- (2) The map $\psi_{0,0}: T_0 \rightarrow T_0$ equals the identity on S .
- (3) The map $\psi_{\infty,\infty}$ equals the identity on T .
- (4) (The semiflow property) $\psi_{ut} \circ \psi_{ts} = \psi_{us}$ for $0 \leq s \leq t \leq u \leq \infty$.

Remarks on Theorem 6.8. Rather than the general F_r trees allowed in [Sko], we are restricting elements $(\psi, S, T) \in \mathcal{M}$ by requiring S, T to be free, simplicial, minimal F_r trees, and by requiring f to satisfy the minimality condition. To see that this is allowable, first note that if \mathcal{M} is defined using ordinary piecewise isometries rather than piecewise isometries satisfying the minimality condition, then Theorem 6.8 is exactly the main result, Theorem 4.8, of [Sko] except for the restrictions to free simplicial trees (and except for correcting the erroneous $\{(s, t) \in [0, 1] \times [0, 1]\}$ occurring in the statement of Theorem 4.8). Then note that for a map $\theta(\psi, s, t) = \psi_{ts}: T_s \rightarrow T_t$ obtained by the constructions of [Sko] or of [GL05], the tree T_t is simplicial by Proposition 3.6 of [GL05], and T_t is free because T is free, and ψ_{ts} satisfies the minimality condition by Propositions 3.4, 3.10, and 3.11 of [GL05].

To complete Step 4, as a consequence of Theorem 6.8 the following chain of compositions is continuous:

$$\mathcal{A} \times [0, \infty] \xrightarrow{(T,s) \mapsto (T,s,\infty)} \mathcal{A} \times \nabla \xrightarrow{(T,s,t) \mapsto (\psi^T, s, t)} \mathcal{M} \times \nabla \xrightarrow{\theta} \mathcal{M} \xrightarrow{\text{Dom}} \mathcal{X}_r$$

This homotopy, which we denote $h: \mathcal{A} \times [0, \infty] \rightarrow \mathcal{X}_r$, has the property that $h(T, 0) = \sigma(T)$ and $h(T, \infty) = T$, in other words h is a homotopy from the map $\sigma: \mathcal{A} \mapsto \mathcal{L}$ to the identity map on \mathcal{A} .

As a consequence of the semiflow property, for each $t \in [0, \infty]$ we obtain a commutative diagram as follows, where we denote $T_t = h(T, t)$:

$$\begin{array}{ccccc} \sigma(T) = T_0 & \xrightarrow{\psi^T} & T_\infty = T & \xrightarrow{k_T^+} & T_+ \\ & \searrow \psi_{t,0}^T & & \nearrow \psi_{\infty,t}^T & \\ & & T_t & & \end{array}$$

We reiterate that the maps $\psi_{t,0}^T, \psi_{\infty,t}^T \in \mathcal{M}$ are edge isometries with the minimality property.

Step 5: We must verify that h takes values in \mathcal{A} , in other words, in the above diagram the tree T_t is a weak train track. For any leaf L of Λ_- , its realization in $\sigma(T)$ denoted $L_{\sigma(T)}$ maps injectively under ψ^T and so also maps injectively under the edge isometry $\psi_{t,0}^T$, with image therefore equal to the realization of L in T_t , denoted L_t . Note that $\psi_{\infty,t}^T$ also restricts to an injection on L_t , and its image equals $\psi^T(L_{\sigma(T)})$, which is the realization of L in T denoted L_T . Since $\psi_{\infty,t}^T$ is an edge isometry it follows that $\psi_{\infty,t}^T$ restricts to an isometry from L_t to L_T . The map $k_T^+ \circ \psi_{\infty,t}^T$ therefore restricts to an isometry from L_t into T_+ . This being true for any leaf L of Λ_- , we have proved that $k_T^+ \circ \psi_{\infty,t}^T: T_t \rightarrow T_+$ is a Λ_- isometry, and so $T_t \in \mathcal{A}$ as desired.

Step 6: We must verify that the homotopy $h: \mathcal{A} \times [0, \infty] \rightarrow \mathcal{A}$ is a proper map, that is, for each compact set $C \subset \mathcal{A}$ the set $h^{-1}(C)$ is compact. This will be a consequence of Theorem 6.4. Since C is compact, there is a closed interval $[a, b] \subset (0, \infty)$ such that $\frac{1}{\text{Len}} \mid C$ takes values in $[a, b]$. For each $T \in \mathcal{A}$, the definition of h and the remark at the end of the verification of Step 2 combine to show that as t goes from 0 to ∞ the value of $\frac{1}{\text{Len}(h(T, t))}$ increases monotonically from $\frac{1}{\text{Len}(h(T, 0))} = \frac{1}{\text{Len}(\sigma(T))} = \frac{\eta}{\lambda L_0 \text{Len}(T)}$ to $\frac{1}{\text{Len}(h(T, \infty))} = \frac{1}{\text{Len}(T)}$. It follows that

$$\frac{1}{\text{Len}(h^{-1}(C))} \subset \left[\frac{\eta}{\lambda L_0} a, b \right]$$

that is, the function $\frac{1}{\text{Len}}$ restricted to $h^{-1}(C)$ takes values in a closed subinterval of $(0, +\infty)$. Since h is continuous, $h^{-1}(C)$ is also closed. Applying Proposition 6.4, it follows that $h^{-1}(C)$ is compact.

Steps 1 through 6 prove that the map $\sigma: \mathcal{A} \rightarrow \mathcal{L}$ is a proper homotopy equivalence. Composing this map with the homeomorphism $\text{Len}: \mathcal{L} \rightarrow (0, +\infty)$, and then composing with the self homeomorphism of $(0, +\infty)$ which is multiplication by the number $\frac{\lambda L_0}{\eta}$, we obtain the map $\text{Len}: \mathcal{A} \rightarrow (0, +\infty)$, proving that this map is a proper homotopy equivalence. This finishes the proof of Theorem 6.1, subject to completing the verification of Step 3(d).

Verifying Step 3(d): Continuity of the map $(\sigma(T), T) \mapsto \psi^T$.

Let $\mathcal{P} = \{(\sigma(T), T) \mid T \in \mathcal{A}\}$, so \mathcal{P} equals the subset of $\mathcal{L} \times \mathcal{A}$ consisting of all (S, T) such that $\text{Len}(T)/\text{Len}(S) = \frac{\eta}{\lambda L_0}$, equivalently, all $(\rho(u), T)$ such that $u \text{Len}(T) = \frac{\eta}{\lambda L_0}$. We shall show that the map $\mathcal{P} \mapsto \mathcal{M}$ taking $(\rho(u), T)$ to ψ^T is continuous. We remind the reader that $\psi^T: \rho(u) \rightarrow T$ is the unique Λ_- isometry that factors the direct limit Λ_- isometry $f_{\rho(u)}^+: \rho(u) \rightarrow T_+$ through the Λ_- isometry $k_T^+: T \rightarrow T_+$. In this argument we shall denote $\psi^T = \psi_u^T$, the superscript being the range and the subscript denoting the domain $\rho(u)$.

Recall from Section 2.1 the Gromov topology on \mathcal{X}_r , with basis element $V(T; \epsilon, t, P)$ determined by $T \in \widehat{\mathcal{X}}_r$, $\epsilon > 0$, a finite subset $t \subset T$, and a finite subset $P \subset F_r$, consisting of all $T' \in \widehat{\mathcal{X}}_r$ for which there exists a finite subset $t' \subset T'$ and a P -equivariant ϵ -almost isometry $\rho \subset t \times t'$ that projects surjectively to t and to t' .

Next we review from [Sko] the Gromov topology on \mathcal{M} . A basis element $U(\psi; \epsilon, s, t, P)$ for this topology is determined for each $\psi: S \rightarrow T$ in \mathcal{M} , $\epsilon > 0$, finite subsets $s \subset S$, $t \subset T$, and finite set $P \subset F_r$. This basis element consists of all $\psi': S' \rightarrow T' \in \mathcal{M}$ for which there exist finite subsets $s' \subset S'$ and $t' \subset T'$, and P -equivariant ϵ -almost isometries $R_0 \subset s \times s'$, $R_1 \subset t \times t'$ such that $(x, y) \in R_0$ implies $(\psi(x), \psi'(y)) \in R_1$. To put this more colloquially, $\psi': S' \rightarrow T'$ is a member of $U(\psi; \epsilon, s, t, P)$ if we can exhibit S' as a member of $V(S; \epsilon, s, P)$ and T' as a member of $V(T; \epsilon, t, P)$ in such a way that the two P -equivariant ϵ -almost isometries that exhibit these memberships are respected by the maps ψ and ψ' . Note the evident continuity of the domain and range maps $\text{Dom}, \text{Ran}: \mathcal{M} \rightarrow \widehat{\mathcal{X}}_r$.

We remark that the Gromov topology on \mathcal{M} as defined in [Sko] uses arbitrary compact subsets of metric spaces, not just finite subsets. However, just as Paulin does in Proposition 4.1 of [Pau89], one obtains a different basis for the same topology on \mathcal{M} by restricting to finite subsets.

The relations that we shall use for ϵ -almost isometries will be restrictions of the Λ_- relations defined in Section 6.2.

Pick $(\rho(u), T) \in \mathcal{P}$, denote $S = \rho(u)$, and let $\psi = \psi_u^T: S \rightarrow T$. Pick a basis element $U(\psi; \epsilon, s, t, P) \subset \mathcal{M}$ about ψ . Corresponding to this are basis elements $V(S; \epsilon, s, P) \subset \mathcal{X}_r$ about S and $V(T; \epsilon, t, P) \subset \widehat{\mathcal{X}}_r$ about T . We must produce P -equivariant ϵ -almost isometries which exhibit membership in these two basis elements.

Applying Lemma 6.3, we obtain a neighborhood $W \subset \mathcal{A}$ of T so that for each $T' \in W$, there is a finite subset $t' \subset T'$ such that if R_1 denotes the restriction to $t \times t'$ of the Λ_- relation between T and T' , then R_1 surjects to t and to t' , and R_1 is a P -equivariant ϵ -almost isometry, thereby exhibiting that $T' \in V(T; \epsilon, t, P)$.

We also need a relation R_0 that exhibits membership in $V(S; \epsilon, s, P)$ as long as $|u - u'|$ is sufficiently small. We will define this relation in two cases, depending on whether $u \leq u'$ or $u' < u$.

When $u \leq u'$ then the semiflow associated to the Stallings fold path \mathcal{L} gives a Λ_- isometry $f_{u',u}: S = \rho(u) \rightarrow \rho(u') = S'$. Set $s' = f_{u',u}(s)$. Regarding $f_{u',u}$ as a subset of $S \times S'$, let $R_0 = f_{u',u} \cap (s \times s')$; to put it another way, R_0 is the restriction of $f_{u',u}$ to s . P -equivariance of R_0 is an immediate consequence of F_r -equivariance of $f_{u',u}$. There is an integer J such that any two points of the finite set s are connected by an edge path with no more than J turns. We choose $|u - u'|$ sufficiently small so that the map $f_{u',u}$ is a fold of length no more than ϵ/J . So, if $x, y \in s$ then the segment between x and y undergoes at most J folds each of length at most ϵ , and so the distance $d_{S'}(f_{u',u}(x), f_{u',u}(y))$ is less than $d_S(x, y)$ by an amount that is no more than $J(\epsilon/J) = \epsilon$. This shows that $f_{u',u} \mid s \times s'$ is an ϵ -almost isometry.

Next assume that $u' \leq u$. The semiflow gives a Λ_- isometry $f_{u,u'}: S' = \rho(u') \rightarrow S = \rho(u)$. Let $s' = f_{u,u'}^{-1}(s)$ and let $R_0 = (f_{u,u'} \mid s')^{-1} \subset s \times s'$. P -equivariance is again immediate. Letting $H(s)$ be the convex hull of s in S , we can think of s' as obtained from s by splittings at certain vertices of $H(s)$. To be precise, for any number $\eta > 0$ such that η is less than half the length of the shortest edge of $H(s)$, if S' is sufficiently close to S in $\widehat{\mathcal{X}}_r$ then s' is obtained from s in the following manner: for certain vertices v in $H(s)$, choose a direction d at v , choose a segment $[v, w]$ representing d of length at most η , partition the directions at v distinct from d into two nonempty subsets $D_1 \cup D_2$, and then split the segment $[v, w]$ into two segments $[v_1, w]$ and $[v_2, w]$, reattach the directions D_1 to v_1 , and reattach the directions D_2 to v_2 . From this construction it follows that there is a constant J depending only on s such that any two points $x, y \in s'$ are connected by an edge path with no more than J turns. This implies that $d_S(f_{u,u'}(x), f_{u,u'}(y))$ is less than $d_{S'}(x, y)$ by an amount that is no more than $J\eta$, and so by choosing $\eta < \epsilon/J$ it follows that $f_{u,u'}^{-1} \mid s \times s'$ is an ϵ -almost isometry.

Finally, we must verify that R_0 and R_1 are respected by the maps ψ_u^T and $\psi_{u'}^{T'}$. That is, we must show that if $(x, y) \in R_0$ then $(\psi_u^T(x), \psi_{u'}^{T'}(y)) \in R_1$. Focussing on the case that $u \leq u'$ (the opposite case $u' \leq u$ being handled similarly), we have $y = f_{u',u}(x)$, and we must verify that

$$(\psi_u^T(x), \psi_{u'}^{T'} \circ f_{u',u}(x)) = (\psi_u^T(x), \psi_u^{T'}(x)) \in R_1 \quad (6.2)$$

R_1 being the restriction of the Λ_- isometry between T and T' , there are two things to check. First, if L is any leaf of Λ_- realized in $\rho(u)$ such that $x \in L$, then since ψ_u^T and $\psi_u^{T'}$ are both Λ_- isometries it follows that $\psi_u^T(x)$ is in the realization of L in T , and $\psi_u^{T'}(x)$ is in the realization of L in T' . Second, we note that

$$k_T^+ \circ \psi_u^T(x) = f_u^+(x) = k_{T'}^+ \circ \psi_u^{T'}(x)$$

completing the verification of (6.2), and thereby completing the verification of Step 3(d) and the proof of Theorem 6.1.

7 Fold lines

In this section we investigate Definition (1) of the axis bundle of a nongeometric, fully irreducible $\phi \in \text{Out}(F_r)$: the union of fold lines converging in $\overline{\mathcal{X}}_r$ to T_-^ϕ in the negative direction and to T_+^ϕ in the positive. Our main result, Theorem 7.2, says that every point on such a fold line is a weak train track for ϕ , thus proving that Definition (1) of the axis bundle implies Definition (2). We also prove the converse, that Definition (2) implies Definition (1); see Proposition 7.1.

Recall from Section 2.3 that a fold path in \mathcal{X}_r is an injective, continuous, proper map $\gamma: I \rightarrow \mathcal{X}_r$ with domain a connected subset $I \subset \mathbf{R}$, for which there is a continuous lift $\hat{\gamma}: I \rightarrow \hat{\mathcal{X}}_r$, marked graphs G_r representing $\hat{\gamma}(r)$ for each $r \in I$, and a family of edge isometries $f_{sr}: G_r \rightarrow G_s$ defined for $s \leq r$ in I and satisfying the semiflow condition $f_{ts} \circ f_{sr} = f_{tr}$. Reparameterizing γ if necessary, we oftentimes impose the following:

Normalization Condition: $\text{Len}(G_r) = e^{-r}$ for each $r \in I$.

If γ is normalized then G_r is uniquely determined up to an isometry that preserves marking, and so the lift $\hat{\gamma}$ is uniquely determined. A translation in the parameter amounts to a change in normalization: for each $T \in \mathbf{R}$, the reparameterization $\gamma'(r') = \gamma(r' + T)$ is also a fold path, and $\hat{\gamma}'(r') = e^T \cdot \hat{\gamma}(r' + T)$.

The set of data $(G_r)_{r \in I}$, $(g_{sr}: G_r \rightarrow G_s)_{r < s \in I}$ is called a *connection* for the fold line γ . While (G_r) is uniquely determined by γ in the presence of the normalization condition, the maps (g_{sr}) are not determined in general, as we saw from the example in Section 5.5.

A normalized fold path γ whose domain I is noncompact only on the positive end then γ is a *fold ray*; and if I is noncompact only on the negative end then γ is called a *split ray*. Denote

$$I_- = \inf(I) \in \mathbf{R} \cup \{-\infty\}, \quad I_+ = \sup(I) \in \mathbf{R} \cup \{+\infty\}$$

As we shall see in the arguments below, if I is noncompact on the negative end then $I_- = -\infty$, in other words $\text{Len}(G_r)$ diverges to $+\infty$ as $r \rightarrow I_-$. But we shall shortly give an example in which I is noncompact on the positive end, γ is indeed proper on the positive end as required, and yet $I_+ \neq +\infty$, that is, $\text{Len}(G_r)$ may converge to a positive number as $r \rightarrow I_+$.

7.1 Examples of fold paths

In this section we give several examples and constructions of fold paths. For each nongeometric, fully irreducible $\phi \in \text{Out}(F_r)$ we construct a periodic fold line through every train track for ϕ , and we construct a fold line for ϕ passing through every weak train track for ϕ ; the latter construction proves that Definition (1) of the axis bundle of ϕ implies Definition (2). In order to motivate some of the later results, we shall also describe a fold ray whose length does not converge to zero, and a split ray which does not converge to a unique point in compactified outer space.

Example: periodic fold lines. Given a nongeometric, fully irreducible $\phi \in \text{Out}(F_r)$ with expansion factor $\lambda = \lambda(\phi)$, we outlined in the introduction the construction of a ϕ -periodic fold line. To give the construction in detail, consider a train track map $g: G \rightarrow G$ representing ϕ . Normalize $G \in \hat{\mathcal{X}}_r$ by requiring $\text{Len}(G) = 1$. Let $G_i = \frac{1}{\lambda^i} G \cdot \phi \in \hat{\mathcal{X}}_r$.

For integers $i \leq j$ let $g_{ji}: G_i \rightarrow G_j$ be the marked homotopy equivalence which agrees with $g^{j-i}: G \rightarrow G$. Following the method described in Section 2.3, choose a Stallings fold sequence for $g_{1,0}: G_0 \rightarrow G_1$, choose an interpolating path for each fold map in that sequence, and concatenate these paths to get a Stallings fold path $\hat{\gamma}^0$ connecting G_0 to G_1 . The path $\hat{\gamma}^0$, when normalized, becomes $\hat{\gamma}^0: [0, \log \lambda] \rightarrow \hat{\mathcal{X}}_r$ with $\hat{\gamma}^0(0) = G_0$ and $\hat{\gamma}^0(\log \lambda) = G_1$. For each integer i define the path $\hat{\gamma}^i: [i \log \lambda, (i+1) \log \lambda] \rightarrow \hat{\mathcal{X}}_r$ to be $\hat{\gamma}^i(t) = (\hat{\gamma}^0 \cdot \phi^i)(t - i \log \lambda)$. Define the path $\hat{\gamma}: \mathbf{R} \rightarrow \hat{\mathcal{X}}_r$ to be the concatenation of $\hat{\gamma}^i$ over all integers i . Connection data for $\hat{\gamma}$ is easily produced starting from the maps g_{ji} and connection data for $\hat{\gamma}^0$. Define the path $\gamma: \mathbf{R} \rightarrow \mathcal{X}_r$ by projectivizing $\hat{\gamma}$. To prove that γ is proper, it suffices to prove that $\gamma(t)$ converges in $\overline{\mathcal{X}}_r$ to T_-^ϕ as $t \rightarrow -\infty$ and to T_+^ϕ as $t \rightarrow +\infty$: by construction the path γ satisfies the ϕ -periodicity condition $\gamma(t - i \log(\lambda)) = \gamma(t) \cdot \phi^i$, and the desired convergence properties follow from the uniform source-sink dynamics of the action of ϕ on compact subsets of \mathcal{X}_r [BFH97].

Example: Weak train tracks. We prove here that Definition (2) of the axis bundle implies Definition (1), by constructing a fold line through every weak train track:

Proposition 7.1. *Let $\phi \in \text{Out}(F_r)$ be fully irreducible and let $\mathcal{A}_\phi \subset \mathcal{X}_r$ denote the set of weak train tracks of ϕ . For each $H \in \mathcal{A}_\phi$ there is a fold line $\gamma: \mathbf{R} \rightarrow \mathcal{X}_r$ passing through H such that $\gamma(t) \rightarrow T_\pm = T_\pm^\phi$ as $t \rightarrow \pm\infty$.*

Proof. Choose a normalization of T_+ . Let $\hat{\mathcal{A}}_\phi \subset \hat{\mathcal{X}}_r$ be the corresponding normalization of \mathcal{A}_ϕ , as discussed in the beginning of Section 6. Denote the normalizations of H and f by $H_0 \in \hat{\mathcal{A}}_\phi$ and by $f_0: \tilde{H}_0 \rightarrow T_+$.

As in the proof of Lemma 2.6, let $h_1: H_0 \rightarrow H_1$ be a fold map whose lift $\tilde{h}_1: \tilde{H}_0 \rightarrow \tilde{H}_1$ factors the map f_0 with induced map $f_1: \tilde{H}_1 \rightarrow T_+$, and such that $(h_1$ and \tilde{h}_1 are complete, meaning that if $e, e' \subset \tilde{H}_0$ are oriented edges being folded by \tilde{h}_1 then the maximal possible initial segments of e and e' that are identified in T_+ are identified immediately in \tilde{H}_1 . Repeat this operation to define an infinite sequence of fold maps

$$H_0 \xrightarrow{h_1} H_1 \xrightarrow{h_2} H_2 \xrightarrow{h_3} \dots$$

with $H_j \in \hat{\mathcal{X}}_r$, and maps $f_i: \tilde{H}_i \rightarrow T_+$, such that $f_{i-1} = f_i \circ \tilde{h}_i$ for integers $i \geq 1$. By induction f_i is a Λ_- isometry and so $H_j \in \hat{\mathcal{A}}_\phi$. By construction, any finite segment of this infinite sequence is a sequence of complete Stallings folds in the sense of Section 2.3. Applying Lemma 2.6 it follows that the combinatorial length of the composition $h_i \circ \dots \circ h_1: H_0 \rightarrow H_i$ goes to infinity as $i \rightarrow +\infty$. Lemma 6.6 therefore implies that the H_i 's eventually leave every compact subset of $\hat{\mathcal{A}}_\phi$. Proposition 6.4 implies that there is a compact fundamental domain for the action of ϕ on \mathcal{A}_ϕ , and the source-sink action of ϕ on compact subsets of \mathcal{X}_r [BFH97] then implies that the closure of \mathcal{A}_ϕ in $\overline{\mathcal{X}}_r$ is $\mathcal{A}_\phi \cup \{T_-, T_+\}$. The sequence H_i therefore accumulates on $\{T_-, T_+\}$ as $i \rightarrow +\infty$. But $\text{Len}(H_i)$ is monotonically decreasing, and so H_i converges to T_+ .

By interpolating the fold maps h_i we produce a fold ray that begins at H and converges to T_+ . By Lemma 6.6 there is Λ_- isometry from some train track G to H . Factoring this as a sequence of folds and interpolating each fold produces a fold ray that begins with G ,

passes through H and converges to T_+ . Finally, preconcatenating with a periodic split ray that ends at G produces the desired fold line γ . \diamond

Example: A fold ray with length $\not\rightarrow 0$. To construct such an example observe that you can do a lot of folds without driving the length to zero, for example by only allowing folds supported in a proper subgraph. For a cheap example, consider any fold ray $\gamma: [0, \infty) \rightarrow \mathcal{X}_r$, with G_s representing $\gamma(s)$. Choose $x \in G_0$. For each $s \geq 0$, let G'_s be a marked graph of rank $r+1$ obtained from G_s by attaching a loop of length 1 based at the point $g_{s0}(x) \in G_s$ a loop of length 1. This yields a fold ray $\rho': [0, \infty) \rightarrow \mathcal{X}_{r+1}$ with $\rho'(s)$ represented by G'_s , but in the positive direction $\text{Len}(G'_s)$ converges to 1, not to zero.

Although we shall prove below that every fold ray in outer space converges to a unique limit point in the boundary of outer space, this example shows that the limiting F_r -tree may not have dense branch points.

Example: A nonconvergent split ray. Let $\rho: (-\infty, 0] \rightarrow \mathcal{X}_r$ be a split ray in rank r outer space \mathcal{X}_r that converges to a tree $S \in \partial\mathcal{X}_r$, and let $\sigma: (-\infty, 0] \rightarrow \mathcal{X}_{r'}$ be a split ray in rank r' outer space $\mathcal{X}_{r'}$ that converges to $T \in \partial\mathcal{X}_{r'}$. Choose connections: a rank r marked graph G_t representing $\rho(t)$ and connection maps $g_{st}: G_t \rightarrow G_s$; and a rank r' marked graph H_t representing $\sigma(t)$ and connection maps $h_{st}: H_t \rightarrow H_s$. Choose base points $x_s \in G_s$ so that $g_{st}(x_t) = x_s$ and $y_t \in H_t$ so that $h_{st}(y_t) = y_s$. Identify $F_r = \pi_1(G_s, x_s)$ and $F_{r'} = \pi_1(H_t, y_t)$. Let $G_s * H_t$ denote the graph obtained from the disjoint union of G_s and H_t by identifying x_s to y_t to form a base point xy_{st} , and so we have identified $\pi_1(G_s * H_t, xy_{st}) = F_{r+r'}$, thus making $G_s * H_t$ into a rank $r+r'$ marked graph. Consider the third quadrant $(-\infty, 0] \times (-\infty, 0]$ in s, t -space. Any continuous path starting from $(s, t) = (0, 0)$ for which $s+t$ decreases strictly monotonically to $-\infty$ defines a split ray in $\mathcal{X}_{r+r'}$, using ρ to split the G factor and σ to split the H factor. Holding $t = 0$ and letting $s \rightarrow -\infty$, we obtain a split ray in $\mathcal{X}_{r+r'}$ converging in $\overline{\mathcal{X}}_{r+r'}$ to a tree \hat{S} on which the $F_{r'}$ factor of $F_{r+r'} = F_r * F_{r'}$ has a fixed point, and so that the F_r action on \hat{S} is isometrically conjugate to S . Similarly, holding $s = 0$ and letting $t \rightarrow -\infty$, we obtain a split ray in $\mathcal{X}_{r+r'}$ converging in $\overline{\mathcal{X}}_{r+r'}$ to a tree \hat{T} on which F_r acts with a fixed point, and $F_{r'}$ acts isometrically conjugate to T . Now construct a split ray as follows: start from $G_0 * H_0$; hold t constant and let s decrease, splitting the G factor, so as to approach very close to \hat{S} ; next, hold s constant and let t decrease, splitting the H factor, so as to approach very close to \hat{T} ; alternate back and forth in this manner, coming ever closer to \hat{S} during the times when G is split, and ever closer to \hat{T} when H is split. The result is a split ray in $\mathcal{X}_{r+r'}$ that accumulates on both \hat{S} and \hat{T} .

7.2 Characterizing fold lines

We have shown in Proposition 7.1 that Definition (2) of the axis bundle of a fully irreducible $\phi \in \text{Out}(F_r)$ implies Definition (1): every weak train track for ϕ lies on some fold line going from T_- to T_+ . The main result of this section is the converse: each point on each fold line going from T_- to T_+ is a weak train track for ϕ , thus proving that Definition (1) implies Definition (2).

Theorem 7.2. *Given a fully irreducible $\phi \in \text{Out}(F_r)$ with attracting and repelling fixed points $T_+, T_- \in \partial\mathcal{X}_r$, if $\gamma: I \rightarrow \mathcal{X}_r$ is a normalized fold line such that*

$$\lim_{t \rightarrow I_+} \gamma(t) = T_+ \quad \text{and} \quad \lim_{t \rightarrow I_-} \gamma(t) = T_- \quad \text{in } \overline{\mathcal{X}}_r$$

then $I = \mathbf{R}$ and for each $t \in \mathbf{R}$, $\gamma(t)$ is a weak train track for ϕ .

Since $\text{Len}(\gamma(t)) = e^{-t}$ it immediately follows that:

Corollary 7.3. *The map $\text{Len}: \gamma(\mathbf{R}) \rightarrow (0, \infty)$ is a homeomorphism.* \diamond

The proof of Theorem 7.2 takes up the rest of Section 7. First we set up the proof, breaking it into three steps, which are carried out in the following subsections.

Choose connection data for γ : for each $t \in I$ a marked graph G_t of length e^{-t} representing $\gamma(t)$, and for each for $s < t \in I$ an edge isometry $g_{ts}: G_s \rightarrow G_t$, so that $g_{ts} \circ g_{sr} = g_{tr}$. Let T_t denote the universal cover of G_t for each $t \in I$. Choose a lift $f_{ts}: T_s \rightarrow T_t$ of each g_{ts} , so that $f_{ts} \circ f_{sr} = f_{tr}$.

We must prove that for each $t \in I$ there is a map $f_t^+: T_t \rightarrow T_+$ which is injective when restricted to each leaf of Λ_- realized in T_t . Along the way we also prove that $I = \mathbf{R}$. There are three steps to the proof.

Step 1: Direct limits of fold rays (Section 7.3). Recall the discussion of direct limits in Section 2.5. Recall also Theorem 2.13 which says that from a train track representative $g: G \rightarrow G$ of ϕ one obtains a sequence of marked graphs $G_i = \frac{1}{\lambda(\phi)^i} G \cdot \phi^i$ whose universal covers fit into a direct sequence with direct limit T_+ . In other words, T_+ is the direct limit of any ϕ -periodic fold ray.

In this section we prove the generalization that for *any* fold ray which converges to T_+ , this convergence is manifested as a direct limit. Regarding the fold ray as a direct system, we first prove that its direct limit is an F_r -tree, and then that T_t converges in $\widehat{\mathcal{X}}_r$ to its direct limit. Using the hypothesis that T_t converges to T_+ , we may identify T_+ with the direct limit. We obtain direct limit maps $f_t^+: T_t \rightarrow T_+$ satisfying the property that $f_t^+ \circ f_{ts} = f_s^+$ for all $s < t$. These maps f_t^+ will be edge isometries. Then, using properties of T_+ , we will prove that $I_+ = +\infty$.

Step 2: Legal laminations of split rays (Section 7.4). Using the fact that T_t converges to T_- in $\overline{\mathcal{X}}_r$ at $t \rightarrow I_-$ we first prove that the injectivity radius of G_t diverges to $+\infty$ as $t \rightarrow I_-$. This immediately implies that $\text{Len}(G_t) \rightarrow +\infty$ and therefore $I_- = -\infty$. Next we focus on the “legal lamination”, a sublamination of the geodesic lamination of F_r , defined as the set of all leaves ℓ of the geodesic lamination such that for all $r < s$ in I , the realization of ℓ in G_r is legal for the map f_{sr} . We prove that the legal lamination is nonempty, and so by the Hausdorff maximal principle has a minimal nonempty sublamination \mathcal{L} . We then prove that this sublamination has zero length in T_- in the sense of [BFH97], and quoting a result of [BFH97] (Lemma 2.19 above) it follows that $\mathcal{L} = \Lambda_-$.

Step 3: Fold lines through weak train tracks (Section 7.5). Putting Steps 1 and 2 together, we prove that the realization of Λ_- in each G_t , when lifted to T_t , is legal for the map f_t^+ , and so G_t is a weak train track.

7.3 Direct limits of fold rays

In this section we carry out Step 1 in the proof of Theorem 7.2, that a fold ray converging to T_+ may be regarded as a direct system with direct limit T_+ .

Consider a normalized fold path $\rho: I \rightarrow \mathcal{X}_r$ which is noncompact in the positive direction. We shall make no other assumptions on ρ until near the end of the argument when we bring in the hypothesis that $\rho(s) \rightarrow T_+$ as $s \rightarrow I_+$.

Choose a normalized connection for ρ : marked graphs G_t representing $\rho(t)$ for each $t \in I$ so that $\text{Len}(G_t) = e^{-t}$; and edge isometries $g_{ts}: G_s \rightarrow G_t$ so that $g_{ts} \circ g_{sr} = g_{tr}$. Denote the universal covers as $T_t = \tilde{G}_t$. We can lift the map g_{ts} to an F_r equivariant map $f_{ts}: T_s \rightarrow T_t$ so that $f_{ts} \circ f_{sr} = f_{tr}$. The maps f_{ts} are all distance nonincreasing edge isometric maps, and so we have a direct system $(T_i)_{i \in I}, (f_{ts}: T_s \rightarrow T_t)_{s < t \in I}$. Let T_∞ be the direct limit and let $f_t^\infty: T_t \rightarrow T_\infty$ be the direct limit maps, distance nonincreasing maps such that $f_s^\infty \circ f_{st} = f_t^\infty$. Since each map f_{ts} restricts to an isometry on each edge of T_s , the map f_s^∞ also restricts to an isometry on each edge. (Despite the notation we do *not* assume at this stage that $I_+ = +\infty$; this will be proved at the appropriate point). The F_r actions on the trees T_t yield an F_r action on T_∞ such that each map f_t^∞ is F_r equivariant: each $\gamma \in F_r$ acts isometrically on each T_s compatible with the maps f_{st} , and the existence clause of the universality property gives the action of γ on T_∞ ; the uniqueness clause gives the action equation $(\gamma \cdot \delta) \cdot x = \gamma \cdot (\delta \cdot x)$ for $\gamma, \delta \in F_r, x \in T_\infty$. Since γ and γ^{-1} are both distance nonincreasing on T_∞ , and since $\gamma\gamma^{-1} = \gamma^{-1}\gamma = \text{Id}$, it follows that γ acts as an isometry on T_∞ .

T_∞ is an \mathbf{R} -tree. First observe that the 0-hyperbolic inequality

$$d(w, x) + d(y, z) \leq \text{Min}\{d(w, z) + d(x, y), d(w, y) + d(x, z)\}$$

holds in each tree T_s , and so evidently holds as well in T_∞ after taking infimums. It follows that T_∞ embeds isometrically in an \mathbf{R} -tree (see e.g. [GdlH91]).

To prove that T_∞ is an \mathbf{R} -tree, it remains to show that if $\xi, \eta \in T_\infty$ and if $a, b > 0$ and $a + b = d(\xi, \eta)$ then there exists $\zeta \in T_\infty$ such that $d(\xi, \zeta) = a$ and $d(\zeta, \eta) = b$. Fix $s_0 \in I$, and choose $x_{s_0}, y_{s_0} \in T_{s_0}$ such that $f_0^\infty(x_0) = \xi$ and $f_0^\infty(y_0) = \eta$. For each $s \geq s_0$ let $x_s = f_{s_0}(x_{s_0})$ and $y_s = f_{s_0}(y_{s_0})$. For all $s_0 \leq s < t \in I$ we have $f_{ts}(x_s) = x_t, f_{ts}(y_s) = y_t, f_s^\infty(x_s) = \xi$ and $f_s^\infty(y_s) = \eta$. For each $s \geq s_0$ choose $v^s \in \overline{x_{s_0}y_{s_0}}$ so that $f_{s,s_0}(v^s) \in \overline{x_sy_s}$ and so that $d(x_s, f_{s,s_0}(v^s)) \geq a$ and $d(f_{s,s_0}(v^s), y_s) \geq b$; this is possible because $\overline{x_sy_s} \subset f_{s,s_0}(\overline{x_{s_0}y_{s_0}})$ and because $d(x_s, y_s) \geq d(\xi, \eta) = a + b$. Some subsequence of the sequence v^s converges to a point $z_{s_0} \in \overline{x_{s_0}y_{s_0}}$. Set $z_s = f_{s,s_0}(z_{s_0})$ and $\zeta = f_{s_0}^\infty(z_{s_0}) = f_s^\infty(z_s)$, for $s \geq s_0$. We shall show that ζ is the desired point.

Observe that for all $s \geq s_0$ we have $d(x_s, z_s) \geq a$. Indeed, if there existed $s \geq s_0$ such that $d(x_s, z_s) < a$ then for all sufficiently large $s' > s$ we would have $d(x_s, f_{s',s_0}(v^{s'})) < a$ which would imply that $d(x_{s'}, f_{s',s_0}(v^{s'})) < a$, a contradiction. By a symmetric argument we have $d(z_s, y_s) \geq b$. Taking infimums as $s \rightarrow \infty$, it follows that $d(\xi, \zeta) \geq a$ and $d(\zeta, \eta) \geq b$.

Next observe that $d(x_s, z_s) + d(z_s, y_s) = d(x_s, y_s)$, and taking infimums as $s \rightarrow \infty$ we obtain $d(\xi, \zeta) + d(\zeta, \eta) = d(\xi, \eta) = a + b$. Combined with the above, it follows that $d(\xi, \zeta) = a$ and $d(\zeta, \eta) = b$.

T_∞ is nontrivial and minimal. Given $s < t \in I \cup \{\infty\}$, and given $x \in T_s$, recall the induced map Df_{ts} taking directions of T_s at x to directions of T_t at $f_{ts}(x)$, and recall that two directions at x are in the same gate relative to f_{ts} if they are identified under Df_{ts} .

Lemma 7.4. *For all sufficiently large s , each point in T_s has ≥ 2 gates relative to f_s^∞ .*

Proof. If d_1 and d_2 are identified by Df_s^∞ then they are identified by Df_{ts} for some finite t . It therefore suffices to prove that if s is sufficiently large then each point in T_s has at least two gates relative to f_{ts} for all finite $t > s$. The property of being in the same gate is invariant under the action of F_r so it suffices to prove that if s is sufficiently large then for each $x \in G_s$ and $t > s$ there are directions d_1 and d_2 at x that are not identified by Dg_{ts} .

Choose $a \in I$. For each $x \in G_a$ consider the following property:

(*) There exists $c \geq a$ such that if $c \leq s \leq t$ then $g_{sa}(x)$ has ≥ 2 gates relative to g_{ts} .

If x is not a vertex of G_a then Property (*) is true with $c = a$, as a consequence of Property (1) in the definition of a fold path. The set of vertices being finite, to complete the proof it remains to verify (*) for each vertex. If x is a vertex, and if there exists $t \geq a$ such that $g_{ta}^{-1}(g_{ta}(x))$ contains a point that is not a vertex of G_a , then Property (*) is true with $c = t$. Consider the set V_a of vertices x of G_a such that for all $t \geq a$, the set $g_{ta}^{-1}(g_{ta}(x))$ consists solely of vertices. As t increases the sets $V_t = g_{ta}(V_a)$ have nonincreasing cardinality, and so for some $c \geq a$ the cardinality of V_t is constant for $t \geq c$. Suppose that $c \leq s \leq t$, $x \in V_a$ and consider $y = g_{ca}(x) \in V_c$. We have shown that $g_{tc}^{-1}(g_{tc}(y)) = \{y\}$. Since g_{tc} maps G_c onto G_t , it maps a neighborhood of y in G_c onto a neighborhood of $g_{tc}(y)$ in G_t , and so there must be at least two directions at y that are not identified by Dg_{tc} . Property (2) in the definition of a fold path implies that the Dg_{sc} images of these two directions are not identified by Dg_{ts} . \diamond

For the remainder of the proof we fix s satisfying Lemma 7.4 and we focus on f_s^∞ . Given $x \in T_s$, a pair of directions at x form a *legal turn* if they are not in the same gate relative to f_s^∞ . A bi-infinite geodesic $\ell \subset T_s$ is *legal* if, for every $x \in \ell$, the turn made by ℓ at x is legal; equivalently, the restriction of f_s^∞ to ℓ is an isometric embedding $\ell \rightarrow T_\infty$. By Lemma 7.4, every direction at every point $x \in T_s$ forms a legal turn with some other direction at x . This implies that every point in T_s lies on some bi-infinite legal geodesic: starting with any edge in T_s , we may extend the edge at each endpoint legally, and continuing to extend by induction we obtain a bi-infinite legal geodesic.

To prove that the action of F_r on T_∞ is nontrivial and minimal, consider $\xi \in T_\infty$, and choose $x \in T_s$ such that $f_s^\infty(x) = \xi$. Let $\ell \subset T_s$ be a bi-infinite legal geodesic passing through x (in this context we do not assume ℓ is a leaf of Λ_-). Let ℓ_-, ℓ_+ be the components of $\ell - x$. Since there are only finitely many F_r -orbits of oriented edges in T_s , the ray ℓ_- must pass over two oriented edges $e \neq e' \subset T_s$ in the same F_r orbit. Let $\gamma_- \in F_r$ take e to e' . Construct a γ_- -invariant line by gluing together translates of the segment of ℓ connecting e to e' . This line must be the axis $A_s(\gamma_-)$ of γ_- in T_s , and the interior of the segment of ℓ from e to e' contains a fundamental domain for $A_s(\gamma_-)$, implying that $A_s(\gamma_-)$ is legal. Doing the same in the other direction, we obtain $\gamma_+ \in F_r$ whose axis $A_s(\gamma_+)$ is legal. It follows that γ_\pm are hyperbolic on T_∞ and their axes $A_\infty(\gamma_\pm)$ are the isometric images under f_s^∞ of the axes $A_s(\gamma_\pm)$; in particular, we now know that the action of F_r on T_∞ is nontrivial. By a standard exercise in group actions on **R**-trees, the axes $A_\infty(\gamma_\pm)$ are contained in the

minimal subtree of T_∞ . By construction, there is a segment of T_s passing through x with endpoints respectively on $A_s(\gamma_\pm)$, and so there is a segment of T_∞ passing through ξ with endpoints respectively on $A_\infty(\gamma_\pm)$, implying that ξ is contained in the minimal subtree of T_∞ . Since $\xi \in T_\infty$ is arbitrary, we proved that the action of F_r on T_∞ is minimal.

Convergence in \mathbf{R}^c . Next we show that T_s converges to T_∞ in \mathbf{R}^c , that is, for each $\gamma \in F_r$ the translation length of γ acting on T_s converges to the translation length of γ acting on T_∞ , as $s \rightarrow I_+$.

Consider the axis $A_s(\gamma) \subset T_s$. For $s < t$ we have $f_{ts}(A_s(\gamma)) \supset A_t(\gamma)$. We may therefore choose fundamental domains $[x_s, y_s] \subset A_s(\gamma)$ for $s \in I$ so that $f_{ts}(x_s) = x_t$, $f_{ts}(y_s) = y_t$. Letting $\xi = f_s^\infty(x_s)$, $\eta = f_s^\infty(y_s)$, it follows that $d(\xi, \eta)$ equals the infimum of the translation lengths. But if $\xi \neq \eta$ then clearly $\cup_{i \in \mathbf{Z}} \gamma^i[\xi, \eta]$ for $i \in \mathbf{Z}$ is an axis for γ with translation distance $d(\xi, \eta)$; whereas if $\xi = \eta$ then $\gamma(\xi) = \eta$ and so γ has translation distance zero.

Remark. So far we have not used the hypothesis that T_s converges to T_+^ϕ . Our arguments therefore give a completely general proof that every fold ray converges in $\overline{\mathcal{X}}_r$ to a unique point $T_\infty \in \partial \mathcal{X}_r$, and that limit may be regarded as a direct limit. We described earlier an example of a fold ray for which length did not converge to zero. For this fold ray it easily follows that limiting F_r -tree T_∞ has nonzero “colength”, meaning that there is a positive lower bound to the length of any finite subtree $\tau \subset T_\infty$ whose translates under F_r cover T_∞ . This implies in turn that T_∞ has a nondegenerate arc whose interior is an open subset of T_∞ , and so the action of F_r on T_∞ is not transitive. The action of F_r on T_∞ therefore has a nontrivial decomposition, in the sense of [Lev94], as a finite graph of actions with dense branch points.

Identification of T_∞ with T_+ . We have shown that T_s converges to T_∞ in \mathbf{R}^c as $s \rightarrow I_+$. Projectivizing, it follows that (the projective class of) T_∞ is a point in $\overline{\mathcal{X}}_r$, and that T_s converges to T_∞ in $\overline{\mathcal{X}}_r$ as $s \rightarrow I_+$. But now we use the hypothesis that T_s converges to T_+ as $s \rightarrow I_+$, to conclude that T_∞ equals T_+ in \mathcal{X}_r , in other words, there exists an F_r -equivariant homothety $T_\infty \approx T_+$. By scaling the metric on T_+ we may conclude that this homothety is an isometry.

We thus obtain edge-isometric morphisms $f_s^+ = f_s^\infty: T_s \rightarrow T_+ = T_\infty$ for $s \in I$, with $f_s^+ = f_t^+ \circ f_{ts}$ for $s < t$ in I .

I_+ equals $+\infty$. From the definition of a normalized fold line combined with the definition of the direct limit, it follows that e^{-I_+} equals the limit of the lengths of the marked graphs $G_s = T_s/F_r$ as $s \rightarrow I_+$. We show that this limit equals zero. If not, then there exists an open subarc $a_s \subset G_s$ for some s , so that for all $t \geq s$, if $x \in a_s$, $y \in G_s$, and $g_{ts}(x) = g_{ts}(y)$ then $x = y$. Letting $A_s \subset T_s$ be a connected lift of a_s , it follows that for all $t \geq s$, if $x \in A_s$, $y \in T_s$, and $f_{ts}(x) = f_{ts}(y)$ then $x = y$. The open arc A_s therefore embeds as an open subarc of T_+ under the direct limit map $f_s^+: T_s \rightarrow T_+$. However, each orbit of the action of F_r on T_+ is dense and so no such open arc exists in T_+ .

This completes Step 1.

7.4 Legal laminations of split rays

In this section we carry out Step 2 in the proof of Theorem 7.2: for any fold ray that converges to T_- , the lamination Λ_- arises naturally as the “legal lamination” of that fold ray. One can think of this as a generalization of Lemma 2.17, in which it was shown that Λ_- is the unique minimal lamination whose leaves are legal with respect to any train track representative of ϕ .

We review the notation: $\rho: I \rightarrow \mathcal{X}_r$ is a normalized fold path assumed to be noncompact in the negative direction; there is a normalized connection for ρ consisting of marked graphs G_t representing $\rho(t)$ for each $t \in I$ so that $\text{Len}(G_t) = e^{-t}$, and edge isometric morphisms $g_{ts}: G_s \rightarrow G_t$ so that $g_{ts} \circ g_{sr} = g_{tr}$. Denote the universal covers as $T_t = \tilde{G}_t$, and lift the map g_{ts} to an F_r equivariant map $f_{ts}: T_s \rightarrow T_t$ of g_{ts} so that $f_{ts} \circ f_{sr} = f_{tr}$.

Injectivity radius. Define the *injectivity radius* of G_t , denoted $\text{inj}(G_t)$, to be the supremum of $\epsilon > 0$ such that for all $x \in G_t$, the ϵ -neighborhood $N_\epsilon(x)$ is a tree and the inclusion map $N_\epsilon(x) \hookrightarrow G_t$ is an isometry with respect to the path length metric on $N_\epsilon(x)$. Note that $\text{inj}(G_t)$ equals half the length of the shortest closed curve on G_t .

For the remainder of this section, choose $\mu_t > 0$ for each $t \in I$ so that

$$\lim_{t \rightarrow I_-} \mu_t T_t = T_- \quad \text{in } \widehat{\mathcal{X}}_r$$

Proposition 7.5. *As $t \rightarrow I_-$ we have: (1) $\text{Len}(\mu_t G_t) \rightarrow 0$; (2) $\mu_t \rightarrow 0$; (3) $\text{inj}(G_t) \rightarrow +\infty$; and (4) $\text{Len}(G_t) \rightarrow +\infty$.*

Proof. Note that the F_r -tree T_- has an arbitrarily short fundamental domain, meaning that for each $\epsilon > 0$ there exists a point $p \in T_-$ and a free basis g_1, \dots, g_r of F_r such that the distances $d_{T_-}(p, f_i \cdot p)$ sum to less than ϵ . This is true on general principles because T_- has dense branch points, and it can be seen very easily in this case by starting with arbitrary p, g_1, \dots, g_r , iterating the action on T_- of any automorphism representing ϕ , and using the fact that ϕ contracts distance on T_- by a factor of $\lambda(\phi^{-1}) > 1$.

Since $\mu_t T_t$ converges to T_- as $t \rightarrow I_-$ in the Gromov–Hausdorff topology on closed subsets of T_- , it follows that for any $\epsilon > 0$, if t is sufficiently close to I_- then there exists $p \in T_t$ and a free basis g_1, \dots, g_r of F_r such that the distances $d_{\mu_t T_t}(p, f_i \cdot p)$ sum to less than ϵ . From this it follows immediately that $\text{Len}(\mu_t G_t) \rightarrow 0$ as $i \rightarrow I_-$. But $\text{Len}(G_s) \geq \text{Len}(G_t)$ for $s < t \in I$, and so $\lim_{t \rightarrow I_-} \text{Len}(G_t)$ exists and is > 0 , possibly $+\infty$. It follows that $\lim_{t \rightarrow I_-} \mu_t = 0$.

Suppose that $\text{inj}(G_t) \not\rightarrow +\infty$ as $t \rightarrow I_-$. Since $\text{inj}(G_t)$ is a nonincreasing function of t , it follows that $\text{inj}(G_t)$ converges to some $R > 0$ as $t \rightarrow I_-$. Recall that \mathcal{C} is the set of nontrivial conjugacy classes in F_r . Let $\mathcal{C}_t = \{C \in \mathcal{C} \mid \text{Len}(C; G_t) \leq 2R\}$, so \mathcal{C}_t is finite and nonempty for all $t \in I$. If $s < t \in I$ then $\mathcal{C}_s \subset \mathcal{C}_t$, because g_{ts} is distance nonincreasing. It follows that \mathcal{C}_t stabilizes for t sufficiently close to I_- , say $t \leq t_0$. Let γ be the minimal length closed curve in G_{t_0} and let $C \in \mathcal{C}_{t_0}$ be the conjugacy class of γ , so $\text{Len}(C; G_t) \leq 2R$ for all $t \in I$. Note that γ is simple and so each representative of C is a generator of a free factor of F_r .

Applying Theorem 5.4 of [BFH97], it follows that $\text{Len}(C; T_-) > 0$. On the other hand,

$$\begin{aligned} \text{Len}(C; T_-) &= \lim_{t \rightarrow I_-} \text{Len}(C; \mu_t T_t) = \lim_{t \rightarrow I_-} \mu_t \text{Len}(C; T_t) \\ &\leq \lim_{t \rightarrow I_-} \mu_t 2R = 0 \end{aligned}$$

a contradiction.

Since $\text{inj}(G_t) \rightarrow +\infty$ as $t \rightarrow I_-$, it follows as well that $\text{Len}(G_t) \rightarrow +\infty$. \diamond

Combining Proposition 7.5 with the normalization condition on the fold ray, we have proved one of the pieces of Theorem 7.2:

Proposition 7.6. $I_- = -\infty$. \diamond

Remark. We have proved that $\lim_{t \rightarrow I_-} \text{Len}(G_t) = +\infty$ via a proof that $\text{inj}(G_t) \rightarrow +\infty$, which itself uses the hypothesis that the split ray converges to T_- . But $\text{inj}(G_t) \not\rightarrow +\infty$ for an arbitrary normalized split ray, as shown by the example described earlier of a split ray in which a certain positive rank subgraph denoted G_0 does not change, as the length of the rest of the subgraph goes to $+\infty$.

Nevertheless, here is a proof that $\lim_{t \rightarrow I_-} \text{Len}(G_t) = +\infty$ which is valid for an arbitrary split ray. If not then by monotonicity we have $\lim_{t \rightarrow I_-} \text{Len}(G_t) = L < +\infty$ and $\text{Len}(G_t) \leq L$ for all $t \in I$. Using this we show for each $C \in \mathcal{C}$ that $\text{Len}(C; G_t)$ an upper bound independent of t , for then by monotonicity the translation length function of G_t has a limit in \mathbf{R}^C whose coordinates are bounded away from zero, and by Lemma 2.4 this corresponds to a point in \mathcal{X}_r , contradicting properness of split rays. To bound $\text{Len}(C; G_t)$, fix $c \in I$, and for $t \leq c \in I$ consider a segment $\alpha \subset T_t$ whose endpoints are identified by f_{ct} . We have $\text{Len}(G_t) - \text{Len}(G_c) \geq \frac{1}{2} \text{Len}(\alpha)$, and so $\text{Len}(\alpha) \leq 2(L - \text{Len}(G_c))$. So if $\gamma \in F_r$ represents C , and if the axis of γ in T_c has K vertices in a fundamental domain, it follows that $\text{Len}(C; G_t) = \text{Len}(\gamma; T_t) \leq \text{Len}(\gamma; T_c) + 2K(L - \text{Len}(G_c))$.

A minimal legal lamination. Consider an element ℓ of the geodesic leaf space $\mathcal{G}F_r$. Let ℓ_t denote the realization of ℓ in T_t . We say that ℓ is a *legal leaf* if for all $s < t \in I$ the restriction of $f_{ts}: T_s \rightarrow T_t$ to the geodesic $\ell_s \subset T_s$ is an embedding, and hence is an isometry onto ℓ_t . The set of legal leaves is a closed, F_r -equivariant subset of $\mathcal{G}F_r$, called the *legal lamination* and denoted $\widetilde{\mathcal{LL}}$. For each $t \in I$ the realization of $\widetilde{\mathcal{LL}}$ in T_t forms an F_r -equivariant sublamination $\widetilde{\mathcal{LL}}_t \subset \mathcal{GL}T_t$ such that f_{ts} pushes $\widetilde{\mathcal{LL}}_s$ forward to $\widetilde{\mathcal{LL}}_t$. Let $\mathcal{LL}_t \subset \mathcal{GL}G_t$ be the downstairs image of $\widetilde{\mathcal{LL}}_t$, and so g_{ts} pushes \mathcal{LL}_s forward to \mathcal{LL}_t .

Fact 7.7. $\widetilde{\mathcal{LL}}$ is not empty.

Proof. Let E_s be the longest edge of G_s and so applying Proposition 7.6 it follows that $\text{Len}(E_s) \geq \frac{1}{3r-3} \text{Len}(G_s) \rightarrow +\infty$ as $s \rightarrow -\infty$. Fixing $t \in I$, for each $s < t$ the path $f_{ts} \mid E_s$ is a legal path in G_t of length equal to $\text{Len}(E_s)$, by item (1) in the definition of foldline. Parameterizing this path by a symmetric interval $[-\text{Len}(E_s)/2, \text{Len}(E_s)/2]$ and extending the domain to \mathbf{R} so that the map is constant on each ray complementary to this interval, we may then pass to a subsequence converging in the compact open topology to a geodesic immersion $\ell_t: \mathbf{R} \rightarrow G_t$. Since ℓ_t is a union of subsegments of paths $f_{ts} \mid E_s$, it follows by

item (1) in the definition of foldline that ℓ_t is legal with respect to all maps $f_{ut}: G_t \rightarrow G_u$, $u \geq t$.

Fixing a sequence $t_1 > t_2 > \dots$ diverging to $-\infty$, by a diagonalization argument we obtain ℓ_{t_i} so that $g_{t_1 t_i}(\ell_{t_i}) = \ell_{t_1}$. Lifting these geodesics to the universal covers T_{t_i} we therefore obtain a legal leaf ℓ . \diamond

We now invoke the Hausdorff maximal principle to conclude that $\widetilde{\mathcal{L}\mathcal{L}}$ has a minimal nonempty, closed, F_r -invariant subset which we denote $\widetilde{\mathcal{L}}$. This is seen more clearly in one of the compact spaces $\mathcal{L}\mathcal{L}_t \subset \mathcal{G}\mathcal{L}G_t$ for any fixed t , because the intersection of any nested sequence of nonempty closed sublaminations of $\mathcal{L}\mathcal{L}_t$ is clearly a nonempty closed sublamination. The realization of $\widetilde{\mathcal{L}}$ in $\mathcal{G}\mathcal{L}T_t$ is denoted $\widetilde{\mathcal{L}}_t$ and the image downstairs in $\mathcal{G}\mathcal{L}G_t$ is denoted \mathcal{L}_t .

Recall that \mathcal{L}_t is said to be *quasiperiodic* if for each compact leaf segment $\alpha \subset \ell$ there exists k such that for each compact leaf segment α' , if $\text{Len}(\alpha') > k$ then α' has a subsegment whose corresponding path in G_t is identical with the path corresponding to α .

We list here the properties of the sublaminations \mathcal{L}_t :

Lemma 7.8. *For each $t \in I$, \mathcal{L}_t is a nonempty, minimal sublamination of $\mathcal{G}\mathcal{L}G_t$, \mathcal{L}_t is quasiperiodic, each leaf of \mathcal{L}_t is aperiodic, and for each $s < t \in I$, g_{ts} maps \mathcal{L}_s to \mathcal{L}_t inducing a bijection of leaves to leaves and taking each leaf to its corresponding leaf by an isometry.*

Proof. For any compact lamination minimality implies quasiperiodicity; for convenience we prove this in the present context. Suppose that \mathcal{L}_t is not quasiperiodic. Fix a leaf segment $\alpha \subset \mathcal{L}_t$ that defies the definition of quasiperiodicity. There is a sequence of leaf segments α'_n with $\text{Len}(\alpha'_n) \rightarrow +\infty$ so that α'_n does not traverse α (we are here confusing a leaf segment with its corresponding path in G_t). Passing to a limit, we obtain a nonempty sublamination of \mathcal{L}_t in which no leaf traverses the segment α , and so this is a sublamination is proper, contradicting minimality.

To prove that no leaf is periodic, if $\ell_t \subset \mathcal{L}_t$ is a periodic leaf then the corresponding leaf $\ell_s \subset \mathcal{L}_s$ is periodic for all $s < t$ and $\text{Len}(\ell_s) = \text{Len}(\ell_t)$. But this implies that $\text{inj}(G_s) \leq \frac{1}{2} \text{Len}(\ell_t)$ for $s < t$, contradicting Proposition 7.5 (3). \diamond

Remark. In an earlier draft we gave a direct but intricate construction which produces sublaminations $\mathcal{L}_t \subset \mathcal{G}\mathcal{L}G_t$ satisfying Lemma 7.8 without invoking the Hausdorff Maximal Principle, by studying the inverse limit as $t \rightarrow -\infty$ of the graphs G_t in the category of compact singular laminations. A careful desingularization process produces the required sublaminations \mathcal{L}_t .

Identifying \mathcal{L} with Λ_- . To complete Step 2, we shall show that \mathcal{L} has length zero in T_- . By applying Lemma 2.19 it will follow that $\mathcal{L} = \Lambda_-$.

We must prove that for any marked graph $G \in \hat{\mathcal{X}}_r$ and every F_r -equivariant morphism $h: \widetilde{G} \rightarrow T_-$ there exists $C \geq 0$ such that for every leaf $\ell \subset \widetilde{G}$ of $\widetilde{\mathcal{L}}_G$ and for any $x, y \in \ell$, $d_{T_-}(hx, hy) \leq C$. To prove this it suffices to check it for *some* $G \in \hat{\mathcal{X}}_r$ and *some* F_r -equivariant morphism $h: \widetilde{G} \rightarrow T_-$; this is because there exists an F_r -equivariant morphism between $\widetilde{G}, \widetilde{G}'$ for any two $G, G' \in \hat{\mathcal{X}}_r$, and because any two F_r -equivariant morphisms

$\tilde{G} \rightarrow T$ are F_r -equivariantly homotopic. We will do this for $G = G_0$ and an arbitrarily chosen F_r -equivariant continuous map $\phi_-: \tilde{G}_0 \rightarrow T_0 \rightarrow T_-$ such that ϕ_- is a homothety when restricted to each edge of T_0 .

Continuing with the notation established earlier in Step 2, denote $T_s^* = \mu_s T_s$, and recall that T_s^* converges to T_- in $\widehat{\mathcal{X}}_r$ as $s \rightarrow -\infty$; when the metric is not important we will use T_s or T_s^* without discrimination. Let \mathcal{L}_s denote \mathcal{L} realized in G_s , and let $\tilde{\mathcal{L}}_s$ denote the realization in T_s or T_s^* . We want to connect the property “ \mathcal{L} has length zero in T_- ” to properties of $\tilde{\mathcal{L}}_s$ for s close to $-\infty$.

We claim that there are F_r -equivariant continuous maps $\phi_s: T_0 \rightarrow T_s^*$ defined for $s \leq 0$ satisfying the following property:

Uniform Length: For any finite geodesic segment $\gamma_0 \subset T_0$, denoting $\gamma_s^\# = (\phi_s)_\#(\gamma_0) \subset T_s^*$ and $\gamma_-^\# = (\phi_-)_\#(\gamma_0) \subset T_-$, we have $\text{Len}(\gamma_s^\#; T_s^*) \rightarrow \text{Len}(\gamma_-^\#; T_-)$ as $s \rightarrow -\infty$.

To prove this claim, fix a base point in T_0 , for $R > 0$ let $B_R \subset T_0$ denote the set of vertices of T_0 that are within distance R of the base point, let G_R denote the set of elements of F_r that move one point of B_R to another, and consider the set $\phi_-(B_R) \subset T_-$. For each $R > 0$ and $\epsilon > 0$, if s is sufficiently large then there exists a G_R -equivariant ϵ -approximation from $\phi_-(B_R)$ to a subset of T_s^* , and by postcomposing ϕ_- with this ϵ -approximation we obtain a G_R -equivariant map $\phi_s^{R, \epsilon}: B_R \mapsto T_s^*$. We may choose functions $R(s) \rightarrow +\infty$ and $\epsilon(s) \rightarrow 0$ so that $\phi_s^{R(s), \epsilon(s)}$ exists for all s sufficiently close to $-\infty$. We may then extend $\phi_s^{R(s), \epsilon(s)}$ in a unique manner to an F_r -equivariant map $\phi_s: T_0 \rightarrow T_s^*$ which is homothetic on each edge. The Uniform Length property for finite geodesic edge paths $\gamma_0 \subset T_0$ follows immediately because the endpoints of γ_0 lie on B_R for sufficiently large R . For an arbitrary geodesic segment of T_0 , which can be written as a finite geodesic edge path with partial edges concatenated at the beginning and at the end, the Uniform Length property follows easily.

As a consequence of the Uniform Length property, observe that there are constants $L > 0$, $s_0 \in I$ such that the map ϕ_- and each of the maps ϕ_s for $s \leq s_0$ are L -lipschitz: just apply the Uniform Length property to one edge in each of the finitely many edge orbits of T_0 . Combining this observation with the bounded cancellation lemma 2.5, there is a constant $B > 0$ such that:

Uniform Bounded Cancellation: The map ϕ_- and each of the maps ϕ_s for $s \leq s_0$ have bounded cancellation constant $\leq B$.

The constants s_0 and B as above shall be fixed for the rest of the argument.

Suppose then that \mathcal{L} does not have length zero in T_- . It follows that there is a finite leaf segment γ_0 of $\tilde{\mathcal{L}}_0$ such that $\text{Len}(\gamma_-^\#; T_-) > 6B$. Applying the Uniform Length property above, for s sufficiently close to $-\infty$ it follows that $\text{Len}(\gamma_s^\#; T_s^*) > 5B$, let's say for all $s \leq s_1$ where $s_1 \leq s_0$. Since \mathcal{L} is minimal there exists a constant K such that for every finite leaf segment β_0 of $\tilde{\mathcal{L}}_0$, if $\text{Len}(\beta_0; T_0) \geq K$ then β_0 contains a translate of γ_0 . By applying Uniform Bounded Cancellation it follows that $\text{Len}(\beta_s^\#; T_s^*) > 3B$ for all $s \leq s_1$ and all leaf segments β_0 of $\tilde{\mathcal{L}}_0$ with $\text{Len}(\beta_0; T_0) \geq K$. We will show that this leads to a contradiction.

Every geodesic segment in G_s of length equal to $\text{inj}(G_s)$ is embedded and passes over at most $2r - 2$ vertices, and therefore intersects some edge in a segment of length at least $C_s = \text{inj}(G_s)/(2r - 3)$. It follows that every geodesic segment of length $C'_s = \text{inj}(G_s) + 2 \text{Len}(G_s)$

contains an entire edge of length at least C_s . There are at most $2r - 1$ edges in G_s and so every geodesic segment of length $C_s'' = [2(2r - 1) + 1]C_s'$ goes at least three times over some edge whose length is at least C_s . As $s \rightarrow -\infty$ we have $\text{inj}(G_s) \rightarrow +\infty$ and so $C_s \rightarrow +\infty$. As $s \rightarrow -\infty$ we also have $\mu_s \text{Len}(G_s) \rightarrow 0$, and since C_s'' is bounded above by a constant times $\text{Len}(G_s)$ it follows that $\mu_s C_s'' \rightarrow 0$ as $s \rightarrow -\infty$.

Let $\tilde{\lambda}_s$ be a leaf of $\tilde{\mathcal{L}}_s$ covering a leaf λ_s of \mathcal{L}_s . There is a subpath $\tilde{\alpha}_s$ of $\tilde{\lambda}_s$ of length at most C_s'' that decomposes as $\tilde{\alpha}_s = \tilde{E}_1 \tilde{\alpha}' \tilde{E}_2 \tilde{\alpha}'' \tilde{E}_3$ where $\tilde{E}_1, \tilde{E}_2, \tilde{E}_3$ are lifts of the same edge E of G_s with length at least C_s . If every crossing of E by the leaf λ_s was immediately followed by $\alpha'' E$ then λ_s would be the periodic leaf $(\alpha'' E)^\infty$. Since λ_s is not periodic it follows that there is a covering translate $\tilde{\lambda}'_s$ of $\tilde{\lambda}_s$ that contains \tilde{E}_2 but not $\tilde{E}_2 \tilde{\alpha}'' \tilde{E}_3$. Similarly, there is a covering translate $\tilde{\lambda}''_s$ that contains \tilde{E}_2 but not $\tilde{E}_1 \tilde{\alpha}' \tilde{E}_2$. Thus the arc $\tilde{\sigma}_s = \tilde{\lambda}_s \cap \tilde{\lambda}'_s \cap \tilde{\lambda}''_s$ satisfies

$$C_s \leq \text{Len}(\tilde{\sigma}_s; T_s) \leq C_s''$$

and so

$$\text{Len}(\tilde{\sigma}_s; T_s^*) \leq \mu_s C_s'' \rightarrow 0 \quad \text{as } s \rightarrow -\infty$$

Let $\beta(s) = f_{0s}(\tilde{\sigma}_s)$. For s sufficiently close to $-\infty$ we have

$$\text{Len}(\beta(s); T_0) = \text{Len}(\tilde{\sigma}_s; T_s) \geq C_s \geq K$$

and so

$$\text{Len}\left((\beta(s))_s^\#; T_s^*\right) > 3B$$

On the other hand, the geodesics $\tilde{\lambda}_s, \tilde{\lambda}'_s, \tilde{\lambda}''_s$, when mapped from T_s to T_0 via f_{0s} , give geodesics in T_0 denoted $\tilde{\lambda}_0, \tilde{\lambda}'_0, \tilde{\lambda}''_0$, and by construction we have

$$\beta(s) \subset \tilde{\lambda}_0 \cap \tilde{\lambda}'_0 \cap \tilde{\lambda}''_0$$

Applying Uniform Bounded Cancellation to the map ϕ_s it follows that if $s \leq s_0$ then $(\beta(s))_s^\# = (\phi_s)_\#(\beta(s))$ is contained in the B -neighborhood of $\tilde{\lambda}_s \cap \tilde{\lambda}'_s \cap \tilde{\lambda}''_s = \tilde{\sigma}_s$, and so for s sufficiently close to $-\infty$ we have

$$\text{Len}\left((\beta(s))_s^\#; T_s^*\right) < 2B + \text{Len}(\tilde{\sigma}_s; T_s^*) < 3B$$

which is a contradiction. This contradiction shows that \mathcal{L} must have length zero in T_- . Since \mathcal{L} is minimal, Lemma 2.19 shows that $\mathcal{L} = \Lambda_-$, completing Step 2.

7.5 Weak train tracks on fold lines

Continuing with the notations of Sections 7.3 and 7.4, we now carry out Step 3, completing the proof of Theorem 7.2.

By Section 7.3 the direct limit of the system of F_r -equivariant maps $f_{st}: T_t \rightarrow T_s$ is T_+ , and we have direct limit maps $f_s: T_s \rightarrow T_+$, each of which is F_r -equivariant and an isometry on each edge. Fix a leaf ℓ of Λ_- and $s \in \mathbf{R}$, and let $\ell_s \subset T_s$ denote the realization of ℓ in T_s . We must show that $f_s|_{\ell_s}$ is an isometry from ℓ_s into T_+ , that is, for each $x, y \in \ell_s$ we must prove $d_{T_+}(f_s(x), f_s(y)) = d_{T_s}(x, y)$. In Section 7.4 we proved that $\Lambda_- = \mathcal{L}$ is the unique minimal legal sublamination. The geodesic ℓ_s is therefore legal, and so $f_{ts}|_{\ell_s}$ is an isometry

for each $t \geq s$. For any $t \geq s$ it follows that $d_{T_t}(f_{ts}(x), f_{ts}(y)) = d_{T_s}(x, y)$. Since the direct sequence $(f_{ts}(x))_{t \geq s}$ determines $f_s(x) \in T_+$, and the sequence $(f_{ts}(y))_{t \geq s}$ determines $f_s(y)$, it follows that $d_{T_+}(f_s(x), f_s(y))$ equals the limit of the constant sequence $d_{T_t}(f_{ts}(x), f_{ts}(y))$ which equals $d_{T_s}(x, y)$, proving that $d_{T_+}(f_s(x), f_s(y)) = d_{T_s}(x, y)$. This proves that each map $f_s: G_s \rightarrow T_+$ is a Λ_- isometry, and so G_s is a weak train track.

References

- [BB97] M. Bestvina and N. Brady, *Morse theory and finiteness properties of groups*, Invent. Math. **129** (1997), no. 3, 445–470.
- [BF] M. Bestvina and M. Feighn, *Outer limits*, preprint.
- [BFH97] M. Bestvina, M. Feighn, and M. Handel, *Laminations, trees, and irreducible automorphisms of free groups*, Geom. Funct. Anal. **7** (1997), 215–244.
- [BFH00] ———, *The Tits alternative for $\text{Out}(F_n)$. I. Dynamics of exponentially-growing automorphisms.*, Ann. of Math. **151** (2000), no. 2, 517–623.
- [BFH04] ———, *Solvable subgroups of $\text{Out}(F_n)$ are virtually Abelian*, Geometriae Dedicata **104** (2004), 71–96.
- [BH92] M. Bestvina and M. Handel, *Train tracks and automorphisms of free groups*, Ann. of Math. **135** (1992), 1–51.
- [Cla05] M. Clay, *Contractibility of deformation spaces of G -trees*, Algebr. Geom. Topol. **5** (2005), 1481–1503 (electronic).
- [CM87] M. Culler and J. W. Morgan, *Group actions on \mathbf{R} -trees*, Proc. London Math. Soc. (3) **55** (1987), no. 3, 571–604.
- [Coo87] D. Cooper, *Automorphisms of free groups have finitely generated fixed point sets*, J. Algebra **111** (1987), no. 2, 453–456.
- [CP94] M. Coornaert and A. Papadopoulos, *Une dichotomie de Hopf pour les flots géodésiques associés aux groupes discrets d’isométries des arbres*, Trans. AMS **343** (1994), no. 2, 883–898.
- [CV86] M. Culler and K. Vogtmann, *Moduli of graphs and automorphisms of free groups*, Invent. Math. **84** (1986), 571–604.
- [FH04] M. Feighn and H. Handel, *Abelian subgroups of $\text{Out}(F_n)$* , preprint, 2004.
- [FLP⁺79] A. Fathi, F. Laudenbach, V. Poenaru, et al., *Travaux de Thurston sur les surfaces*, Astérisque, vol. 66–67, Société Mathématique de France, 1979.
- [GdlH91] E. Ghys and P. de la Harpe, *Infinite groups as geometric objects (after Gromov)*, Ergodic theory, symbolic dynamics and hyperbolic spaces (Trieste, 1989) (T. Bedford, M. Keane, and C. Series, eds.), Oxford Univ. Press, 1991, pp. 299–314.
- [GJLL98] D. Gaboriau, A. Jaeger, G. Levitt, and M. Lustig, *An index for counting fixed points of automorphisms of free groups*, Duke Math. J. **93** (1998), no. 3, 425–452.
- [GL05] V. Guirardel and G. Levitt, *The outer space of a free product*, preprint, arXiv:math.GR/0501288, 2005.
- [HM06a] M. Handel and L. Mosher, *The expansion factors of an outer automorphism and its inverse*, Trans. AMS (2006), To appear. Preprint, arXiv:math.GR/0410015.

- [HM06b] ———, *Parageometric outer automorphisms of free groups*, Trans. AMS (2006), To appear. Preprint, arXiv:math.GR/0410018.
- [HT85] M. Handel and W. Thurston, *New proofs of some results of Nielsen*, Adv. in Math. **56** (1985), no. 2, 173–191.
- [Lev94] G. Levitt, *Graphs of actions on \mathbf{R} -trees*, Comment. Math. Helv. **69** (1994), no. 1, 28–38.
- [LL00a] G. Levitt and M. Lustig, *Most automorphisms of a hyperbolic group have very simple dynamics*, Ann. Sci. École Norm. Sup. (4) **33** (2000), no. 4, 507–517.
- [LL00b] ———, *Periodic ends, growth rates, Hölder dynamics for automorphisms of free groups*, Comment. Math. Helv. **75** (2000), no. 3, 415–429.
- [LL03] ———, *Irreducible automorphisms of F_n have north-south dynamics on compactified outer space*, J. Inst. Math. Jussieu **2** (2003), no. 1, 59–72.
- [LL04] J. Los and M. Lustig, *The set of train track representatives of an irreducible free group automorphism is contractible*, preprint, December 2004.
- [Lus99] M. Lustig, *Structure and conjugacy for automorphisms of free groups*, preprint, 1999.
- [McC85] J. McCarthy, *A “Tits-alternative” for subgroups of surface mapping class groups*, Trans. AMS **291** (1985), no. 2, 582–612.
- [MS93] H. Masur and J. Smillie, *Quadratic differentials with prescribed singularities and pseudo-Anosov diffeomorphisms*, Comment. Math. Helv. **68** (1993), no. 2, 289–307.
- [Nie86] J. Nielsen, *Investigations in the Topology of Closed Orientable Surfaces, I*, Jakob Nielsen: Collected Mathematical Papers, Volume 1, Birkhäuser, 1986, Translation by John Stillwell of: Untersuchungen zur Topologie der geschlossenen zweiseitigen Flächen. Acta. Math., 50 (1927), 189–358.
- [Pau89] F. Paulin, *The Gromov topology on \mathbf{R} -trees*, Topology Appl. **32** (1989), no. 3, 197–221.
- [Sko] R. Skora, *Deformations of length functions in groups*, preprint.
- [Sko90] ———, *Splittings of surfaces*, Bull. AMS **23** (1990), 85–90.
- [Spa81] E. Spanier, *Algebraic topology*, Springer, New York—Berlin, 1981.
- [Sta83] J. Stallings, *Topology of finite graphs*, Inv. Math. **71** (1983), 551–565.
- [Vog02] K. Vogtmann, *Automorphisms of free groups and outer space*, Proceedings of the Conference on Geometric and Combinatorial Group Theory, Part I (Haifa, 2000), vol. 94, 2002, pp. 1–31.
- [Whi93] T. White, *Fixed points of finite groups of free group automorphisms*, Proc. AMS **118** (1993), no. 3, 681–688.